Time–Entropy Duality and the Lax Double Category of Thermodynamics

I. Stochastic and Quantum Markov Dynamics

June 2026

Abstract

We develop a unified framework for the structural relationship between time and thermodynamic entropy in stochastic thermodynamics and quantum Markov dynamics. The central claim is the *Duality Proposition*: time and entropy are not independent fundamental quantities, nor is one reducible to the other; both arise as dual quantities derived from the symmetry structure of a path-space action functional. This duality is made mathematically precise through a *lax double category* $\mathbb{T}\mathrm{herm}$ whose horizontal morphisms correspond to reversible (time-like) dynamics and whose vertical morphisms correspond to irreversible (dissipative) processes.

The non-trivial structure of comparison cells—arising from the non-commutativity $[H_A, H_S] \neq 0$ of the symmetric and antisymmetric generators in the GNS (Gelfand–Naimark–Segal) representation—is identified as the *categorical mixing cost* $\mathcal{A}(\chi)$: a geometric measure of the coupling between the two dynamical directions. We carefully distinguish this categorical cost from the housekeeping entropy production σ_{hk} of stochastic thermodynamics. The two quantities have nested (not equal) zero sets, $\{H_A = 0\} \subsetneq \{[H_A, H_S] = 0\}$, and on the generic locus they vanish at different orders in the cycle affinity ε: $\mathcal{A}(\chi) \sim \varepsilon$ while $\sigma_{\mathrm{hk}} \sim \varepsilon^2$. Near equilibrium $\mathcal{A}(\chi)$ therefore dominates and is the more fundamental of the two.

The framework is established through four interconnected results: (1) the Duality Proposition, grounded in Noether-charge thermodynamics (Sasa–Yokokura) and its extension to gradient flows and GENERIC (Beyen–Maes); (2) a category-theoretic naturality analysis valid for arbitrary probability spaces: the time-like and entropy-like directions are the isotypic components of the path-reversal involution $\theta(L) = -L^{\dagger_\pi}$, characterized by an ambidextrous adjunction, unique among all natural splittings, and globally jointly exhaustive via the Cartan (polar) factorization; (3) explicit verification in finite-state Markov chains via the GNS decomposition $H = H_S + H_A$, with complete proofs of the structural lemmas, *exact all-order interchange cocycle identities* (together with two-sided finite-size bounds for the comparison-cell cost), the Order Separation Proposition, an exact closed form for the two-bath model, and a channel-resolved lift restoring $H_A \neq 0$ for multi-bath two-state systems, together with computational confirmation of the universal scaling exponents across eight independent chains; and (4) a rigorous connection to quantum thermodynamic uncertainty relations: the structural framework transfers verbatim to Lindblad dynamics via a KMS-superoperator involution, quantum coherence corrections $C(\tau)$ are identified as quantum comparison cells, and quantum order separation is proved—with quantum detailed balance retaining the coherent

rotation in the reversible direction, so that coherence shifts the separation exponents from $(1, 2)$ to $(0, 2)$.

The relativistic extension of this framework—observer-indexed prestacks, the Unruh effect, Bekenstein–Hawking entropy, and logarithmic corrections—is developed in the companion paper (Part II).

Note on reproducibility. All numerical results in Section 6 have been independently confirmed by computer; details are available from the authors.

Contents

1 Introduction **4**

2 Four Recent Advances in Entropy–Time Relations **5**
2.1 The Landauer limit . . . 5
2.2 Entropy as a Noether invariant: Sasa–Yokokura . . . 5
2.3 Thermodynamic cost of timekeeping: Erker–Pearson . . . 5
2.4 Extension to gradient flows and GENERIC: Beyen–Maes . . . 5

3 Against Strong Reductionism **5**
3.1 The evidence for strong reductionism . . . 5
3.2 Four critical difficulties . . . 5

4 The Duality Proposition **6**
4.1 Formulation . . . 6
4.2 Three corollaries . . . 6
4.3 Resolution of the four difficulties . . . 6

5 Category-Theoretic Naturality of the Duality **7**
5.1 The naturality problem . . . 7
5.2 First and second attempts (rejected) . . . 7
5.3 The double category $\mathbb{T}\mathrm{herm}$ . . . 7
5.4 Axioms: lax double category . . . 7
5.5 Universal characterization of the two directions . . . 8
5.6 The universal property in full generality: time reversal as an involution . . . 9
5.7 Physical predictions . . . 12

6 Verification in Finite-State Markov Chains **12**
6.1 Setup and basic quantities . . . 12
6.2 GNS decomposition: the two generators . . . 13
6.3 Failure of the naive Helmholtz split . . . 14
6.4 Axiom verification . . . 14
6.5 Lax interchange and the comparison cell . . . 15
6.6 Exact coherence: the interchange cocycle . . . 16
6.7 Categorical mixing cost vs. housekeeping entropy production . . . 19
6.8 Zero-set structure . . . 21
6.9 Two-bath model: exact closed form . . . 22
6.10 Channel-resolved lift: restoring $H_A \neq 0$ for the two-bath system . . . 24

7 From the Path Action to the GNS Split **25**

8 Quantum Thermodynamic Uncertainty Relations **26**
8.1 Quantum Fisher information and commutator representation . . . 27
8.2 Hasegawa's quantum TUR . . . 27
8.3 The superoperator involution and the transfer of the structural framework . . . 27
8.4 Quantum order separation . . . 29
8.5 Main theorem: unified TUR . . . 30
8.6 Secular error propagation: quantifying claim (ii) . . . 31

8.7 Coherent–incoherent correspondence functor Φ . . . 32
8.8 Explicit calculation: field-driven two-level system . . . 32
8.9 New predictions . . . 32

9 Discussion and Outlook **32**
9.1 What has been achieved . . . 32
9.2 Resolution record . . . 33
9.3 Remaining open problems . . . 34

10 Conclusion **35**

A GNS Transformation and the D1–D2 Construction **37**
A.1 The ψ-space transformation . . . 37
A.2 Proof of $H_A\sqrt{\pi} = H_S\sqrt{\pi} = 0$ . . . 37
A.3 D1–D2 construction . . . 37
A.4 Computational verification . . . 37
A.5 Relation to the Helmholtz formulation . . . 37

1 Introduction

The relationship between time and thermodynamic entropy is among the deepest problems in physics. Since Boltzmann's H-theorem, the second law of thermodynamics has been understood to define a preferred direction in time. Recent advances in stochastic thermodynamics [4, 6] and quantum information have elevated this relationship to a new level of scrutiny.

Four independent recent developments motivate the present work:

1. The experimental verification of the Landauer limit [3].
2. The characterization of thermodynamic entropy as a Noether invariant by Sasa and Yokokura [4].
3. The experimental confirmation by Pearson et al. [6] that clock accuracy is linearly bounded by entropy production: $N \lesssim \Delta S_{\text{tick}}/k_B$.
4. The extension of the Noether framework to gradient flows and GENERIC by Beyen and Maes [7, 8].

The Duality Proposition.

Duality Proposition

Time and entropy are *dual quantities* derived from the same underlying structure—the symmetry group of a path-space action functional $\mathcal{A}[z(\cdot)]$. The time parameter t arises as the parameter associated with exact time-translation symmetry; the entropy S arises as the Noether charge of a deformed ("entropic") symmetry. Neither is more fundamental; both are faces of the same action-functional geometry.

Structure of the paper. Section 2 reviews the four key advances. Section 3 analyzes strong reductionism and its difficulties. Section 4 develops the Duality Proposition. Section 5 provides the category-theoretic naturality analysis and defines the lax double category $\mathbb{T}\text{herm}$. Section 6 verifies the structure in finite-state Markov chains via the GNS decomposition, with complete proofs of all structural results and the Order Separation Proposition. Section 7 traces the conceptual route from the path-space action to the GNS split. Section 8 connects the framework to quantum thermodynamic uncertainty relations. Section 9 discusses what has been achieved and lists open problems. Appendix A gives the GNS construction in full.

Relation to the companion paper. The extension of this framework to relativistic systems—special-relativistic observer-indexed prestacks, the Klein–Gordon field with BDNK dissipation, the Unruh effect, Schwarzschild spacetime and Bekenstein–Hawking entropy, and the higher-categorical treatment of logarithmic corrections—is a thematically distinct development and is presented separately in Part II [26]. The present paper is self-contained and does not depend on Part II.

2 Four Recent Advances in Entropy–Time Relations

2.1 The Landauer limit

Landauer (1961) showed that the irreversible erasure of one bit dissipates at least $k_B T \ln 2$ of heat [1]. Bérut et al. [3] verified this experimentally. Reversible computation carries no thermodynamic cost (Bennett 1973 [2]). The lesson is that *irreversibility*, not computation per se, is what carries an entropic price—a first hint that the reversible/irreversible distinction is structural rather than incidental.

2.2 Entropy as a Noether invariant: Sasa–Yokokura

Sasa and Yokokura [4] showed that the non-uniform time translation

$$t \mapsto t + \eta\, \hbar\beta(E, \alpha) \tag{1}$$

constitutes a symmetry of the thermodynamic action (restricted to quasi-static trajectories), and that the associated Noether charge is precisely the thermodynamic entropy S. Entropy thus appears not as a measure of ignorance but as the conserved charge of a *deformed time translation*.

2.3 Thermodynamic cost of timekeeping: Erker–Pearson

Erker et al. [5] argued that any autonomous clock producing N distinguishable ticks must dissipate, with accuracy linearly bounded by entropy production, $N \lesssim \Delta S_{\text{tick}}/k_B$; Pearson et al. [6] verified this experimentally. Measuring time *costs* entropy: the two quantities are operationally interlocked.

2.4 Extension to gradient flows and GENERIC: Beyen–Maes

Beyen and Maes [7, 8] extended the Noether framework to GENERIC systems

$$\dot{z} = L\nabla E + M\nabla S, \qquad L\nabla S = 0, \quad M\nabla E = 0, \tag{2}$$

in which a Hamiltonian (reversible) generator L and a dissipative (metric) generator M coexist with mutual degeneracy conditions. The GENERIC split is the finite-dimensional ancestor of the two-generator structure that the present paper develops categorically.

3 Against Strong Reductionism

3.1 The evidence for strong reductionism

The structural parallels above—together with holographic entanglement entropy (Ryu–Takayanagi [12]) and the thermodynamic derivation of the Einstein equations (Jacobson [11])—suggest the strong thesis that entropy *is* time: that temporal structure is nothing but coarse-grained entropic structure.

3.2 Four critical difficulties

3.2.1 Metric arbitrariness. $dS \geq 0$ defines an arrow but not a metric. Different clocks yield process-dependent "seconds"; entropy alone cannot calibrate duration.

3.2.2 Disappearance of time in the reversible limit. If time were constituted by entropy increase, $dS = 0$ would render quasi-static processes "timeless"—yet quasi-static processes take time, indeed infinite time.

3.2.3 Incompatibility with relativistic spacetime. An observer-dependent entropy introduces a second observer-dependence over and above Lorentz covariance, with no canonical way to reconcile the two (a difficulty open since Connes–Rovelli 1994 [10]). This difficulty is the point of departure for the companion paper [26].

3.2.4 Inconsistency with unitary foundations. The Schrödinger equation is unitary; placing entropy at the foundational level requires coarse-graining as a primitive, which reverses the usual order of explanation.

4 The Duality Proposition

4.1 Formulation

Duality Proposition (Precise Form)

Let the dynamics of a physical system be governed by a path-space action $\mathcal{A}[z(\cdot)]$. Time t appears as the parameter of exact time-translation symmetry; entropy S appears as the Noether charge of a deformed symmetry. Both arise from distinct symmetry directions of the same $\mathcal{A}$. Neither is primary.

4.2 Three corollaries

4.2.1 Geometrization of the second law. $dS \geq 0$ follows from the non-negativity of the large-deviation rate functional ($\mathcal{I}[z] \geq 0$): the second law is the statement that the realized path minimizes the action among paths with the same endpoints.

4.2.2 Unified origin of thermodynamic inequalities. $\Delta S_{\min}(\mathcal{T}) = \inf\{\mathcal{I}[z] : z \in \mathcal{C}_{\mathcal{T}}\}$: every process-class $\mathcal{T}$ carries a minimal entropic cost given by a constrained variational problem on the same action.

4.2.3 Emergent time and standard time as two faces. External time arises when time-translation symmetry is manifest; entropic time arises when only the modular flow (KMS structure) makes the symmetry manifest.

4.3 Resolution of the four difficulties

Metric arbitrariness: the action functional supplies the normalization that bare entropy increase lacks. *Reversible limit*: $dS = 0$ means the entropic symmetry is unbroken, but the time-translation symmetry—and hence time—persists. *Relativistic*: addressed by the observer-indexed stack construction of Part II [26]. *Unitary foundations*: entropy production arises from selective breaking of the deformed symmetry, not from a primitive coarse-graining.

5 Category-Theoretic Naturality of the Duality

5.1 The naturality problem

When a path-space action admits infinitely many symmetries, why do the time-translation generator G_t and the entropic generator G_S occupy a privileged position? Without an answer, the Duality Proposition risks being a relabeling rather than an explanation.

5.2 First and second attempts (rejected)

A naive functor $\Sigma : \mathbf{Sys} \to \mathbf{Sym}$ assigning to each system its symmetry groupoid cannot *select* two distinguished symmetries without assuming the time/entropy correspondence it is meant to explain—a circularity. A direct adjoint pair Time $\dashv$ Entropy between a category of time-ordered processes and a category of entropy-monotone flows restricts the admissible morphisms so severely that most physical systems fall outside the framework. The correct setting keeps *both* directions inside a single two-dimensional structure.

5.3 The double category $\mathbb{T}\mathrm{herm}$

Definition 5.1: The double category $\mathbb{T}\mathrm{herm}$

- **Objects**: thermodynamic states (M, ρ)—a state space M with reference (stationary) measure π and a probability distribution ρ on M.
- **Horizontal morphisms**: reversible maps that are $L^2(\pi)$-isometries, preserving the χ^2-divergence: $D_{\chi^2}(h \cdot \rho \,\|\, \pi) = D_{\chi^2}(\rho \,\|\, \pi)$. In the GNS ψ-space representation these are orthogonal transformations generated by $-H_A$ (Section 6.2 and Appendix A).
- **Vertical morphisms**: stochastic kernels contracting relative entropy: $D_{\mathrm{KL}}(v \cdot \rho \,\|\, \pi) \leq D_{\mathrm{KL}}(\rho \,\|\, \pi)$. Generated by $-H_S$.
- **Two-cells**: compatibility squares with cost $\mathcal{A}(\square) \in [0, \infty]$.

Remark 5.1 (Why divergences, not Shannon entropy)**.** *One might be tempted to define horizontal morphisms as Shannon-entropy-preserving and vertical morphisms as Shannon-entropy-increasing. Both conditions fail for the generators at hand: the flow generated by the antisymmetric part does not preserve the probability simplex in the naive formulation (Theorem 6.1 below), and the dissipative flow* decreases *Shannon entropy whenever the stationary distribution is less uniform than the initial one (for $k_+ = 100$, $k_- = 1$ the stationary entropy $S(\pi) \approx 0.056$ nats lies far below $S(1/2, 1/2) = \ln 2$). The divergence-based conditions of Definition 5.1 are the correct invariants: $D_{\chi^2}(\cdot\|\pi)$-preservation for the reversible direction and $D_{\mathrm{KL}}(\cdot\|\pi)$-contraction (data-processing inequality) for the dissipative one.*

5.4 Axioms: lax double category

We impose Ehresmann-style axioms [14], with the interchange axiom in *lax* form:

- **(D1)** Horizontal and vertical morphisms are well-defined and closed under identity.

- **(D2)** Horizontal composition: $h_{t_1} \circ h_{t_2} = h_{t_1+t_2}$ forms a one-parameter group of D_{χ^2}-preserving maps.
- **(D3)** Vertical composition: $v_{s_1} \circ v_{s_2} = v_{s_1+s_2}$ forms a one-parameter semigroup of D_{KL}-contractive kernels.
- **(D4)** Two-cell composition is associative.
- **(D5′) (Lax interchange)**: there exists a *non-invertible* 2-cell $\chi_{t,s}\colon h_t\, v_s \Rightarrow v_s\, h_t$ with leading-order cost

$$\mathcal{A}(\chi) = \tfrac{1}{2}\left\|[H_A, H_S]\right\|_F \cdot t\, s, \tag{3}$$

satisfying lax bilinearity (the pentagon condition at order ts) and the lax Jacobi/hexagon condition at order $(t^2 s, ts^2)$; both are proved in Section 6.5. In fact, exact all-order coherence identities hold (Theorem 6.2); the pentagon and hexagon are their linearizations.

Why lax, not pseudo. Grandis–Paré [13] define a *pseudo*-double category as one in which the comparison cells are *invertible* globular 2-isomorphisms, and in which it is the associativity/unitality of one composition direction—not the interchange law—that is weakened. The cell χ of (D5′) fails both criteria: its cost $\mathcal{A}(\chi) > 0$ is a genuine irreversible defect whenever $[H_A, H_S] \neq 0$, so no inverse cell exists; and what fails in $\mathbb{T}\mathrm{herm}$ is precisely interchange. A structure with strict compositions in both directions but only a lax (non-invertible) interchange constraint is a **lax double category**, and that is the correct description of $\mathbb{T}\mathrm{herm}$. The full lax coherence tower is established in exact (all-orders) form by the interchange cocycle identities of Theorem 6.2.

5.5 Universal characterization of the two directions

Central Proposition (proved form). The horizontal and vertical directions of $\mathbb{T}\mathrm{herm}$ are the ± 1 isotypic components of the *path-reversal involution* $\theta(L) = -L^{\dagger_\pi}$. Each is characterized by a universal property—an ambidextrous adjunction (Theorem 5.1)—and the splitting is the *unique* natural one (Theorem 5.2); globally, the two directions jointly exhaust all invertible dynamics via the Cartan factorization (Theorem 5.3). No entropy functional enters at any point.

The non-circularity of this characterization is established by the following decomposition result, proved in the Markov-chain setting in Section 6.2: the generator of any finite-state Markov chain splits uniquely into a part antisymmetric and a part symmetric with respect to the stationary geometry, *and the construction of the antisymmetric part makes no reference to any entropy functional*—it is determined by the $L^2(\pi)$ inner product alone.

Lemma 5.1 (Existence of the split)**.** *For every finite-state generator L with stationary distribution π there exist A antisymmetric and B symmetric in $L^2(\pi)$ with $L = A + B$.*

Proof. Take $A = \frac{1}{2}(L - L^\dagger_\pi)$ and $B = \frac{1}{2}(L + L^\dagger_\pi)$, where $L^\dagger_\pi$ is the $L^2(\pi)$-adjoint. Antisymmetry and symmetry follow from $(L^\dagger_\pi)^\dagger_\pi = L$, and $A + B = L$ by construction. No entropy functional appears. □

Lemma 5.2 (Uniqueness of the split)**.** *The decomposition of Lemma 5.1 is unique.*

Proof. If $L = A' + B'$ is another such decomposition, then $D := A' - A = B - B'$ is simultaneously antisymmetric and symmetric in $L^2(\pi)$, so $\langle Df,\, f\rangle_\pi = -\langle f,\, Df\rangle_\pi = -\langle Df,\, f\rangle_\pi$, giving $\langle Df,\, f\rangle_\pi = 0$ for all f. Polarizing, $\langle Df,\, g\rangle_\pi + \langle Dg,\, f\rangle_\pi = 0$; combined with symmetry of D this gives $\langle Df,\, g\rangle_\pi = 0$ for all f, g. Since $\pi_i > 0$ for all i, the form $\langle \cdot,\, \cdot\rangle_\pi$ is positive-definite, hence non-degenerate, and $D = 0$. □

These lemmas treat a fixed finite-state chain. The next subsection upgrades them to the general statement—arbitrary probability spaces, bounded generators, no dimension restriction—by isolating the structure behind them: time reversal is an *involution*, and the two directions are its isotypic components. The loose phrasing "G_t is left adjoint to the forgetful functor"—easily asserted, but meaningless without a precise definition of the categories involved—is thereby replaced by proved universal properties.

5.6 The universal property in full generality: time reversal as an involution

Throughout this subsection, (X, Σ, π) is an arbitrary probability space, $\mathcal{H} = L^2(\pi)$ the associated real Hilbert space, and generators are bounded operators $L \in B(\mathcal{H})$ (for Markov chains this means uniformly bounded jump rates; the unbounded case is discussed in Remark 5.3). We write $L^{\dagger_\pi}$ for the adjoint with respect to $\langle \cdot, \cdot\rangle_\pi$.

Definition 5.2: Path-reversal involution

Define $\theta\colon B(\mathcal{H}) \to B(\mathcal{H})$ by

$$\theta(L) \;:=\; -L^{\dagger_\pi}. \tag{4}$$

Physically, $L^{\dagger_\pi}$ generates the time-reversed process (the π-reversal), and the sign reverses the flow direction, so $\theta(L)$ is the generator of the time-reversed dynamics run backwards; $\theta(L) = L$ says exactly that reversing the process and running it backwards reproduces the original flow—reversibility.

Lemma 5.3 (Cartan symmetric-pair structure)**.** *θ is an involutive Lie-algebra automorphism of $\mathfrak{g} := B(\mathcal{H})$: $\theta^2 = \mathrm{id}$ and $\theta([L, M]) = [\theta L, \theta M]$. Its eigenspaces*

$$\mathfrak{k} := \{L :\, L^{\dagger_\pi} = -L\} \quad (\theta = +1), \qquad \mathfrak{p} := \{L :\, L^{\dagger_\pi} = L\} \quad (\theta = -1) \tag{5}$$

give $\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}$ with the Cartan relations

$$[\mathfrak{k}, \mathfrak{k}] \subseteq \mathfrak{k}, \qquad [\mathfrak{k}, \mathfrak{p}] \subseteq \mathfrak{p}, \qquad [\mathfrak{p}, \mathfrak{p}] \subseteq \mathfrak{k}. \tag{6}$$

Here $\mathfrak{k}$ consists of the generators of the orthogonal group of $(\mathcal{H}, \langle \cdot, \cdot\rangle_\pi)$—the horizontal direction—and $\mathfrak{p}$ of the π-symmetric (dissipative) generators—the vertical direction. This is the standard structure of a Cartan symmetric pair [16]. In the GNS ψ-space picture, the projections onto $\mathfrak{k}, \mathfrak{p}$ are exactly $L \mapsto H_A$, $L \mapsto H_S$ of (16).

Proof. $\theta^2(L) = -(-L^{\dagger_\pi})^{\dagger_\pi} = L$. For the automorphism property, using that $\dagger_\pi$ is an anti-homomorphism,

$$[\theta L, \theta M] = [L^{\dagger_\pi}, M^{\dagger_\pi}] = L^{\dagger_\pi} M^{\dagger_\pi} - M^{\dagger_\pi} L^{\dagger_\pi} = (ML)^{\dagger_\pi} - (LM)^{\dagger_\pi} = -([L, M])^{\dagger_\pi} = \theta([L, M]).$$

Since θ is multiplicative on brackets, the bracket of θ-homogeneous elements is homogeneous with eigenvalue the product of the eigenvalues; this is (6). The eigenspace identification and the equality $\frac{1}{2}(1+\theta)L = \frac{1}{2}(L - L^{\dagger_\pi})$ are immediate. □

Corollary 5.1 ($\mathbb{Z}/2$-grading of the comparison-cell hierarchy)**.** *Every nested commutator in the letters $H_A \in \mathfrak{k}$, $H_S \in \mathfrak{p}$ is θ-homogeneous, of type $\mathfrak{p}$ (dissipative) if it contains an odd number of H_S-letters and of type $\mathfrak{k}$ (rotational) if even. In particular the comparison 2-cell $[H_A, H_S]$ is* dissipative-type*; among the 3-cells, $[H_A, [H_A, H_S]]$ is dissipative-type while $[H_S, [H_A, H_S]]$ is rotational-type: two dissipative letters interfere to produce a rotation. Every cell of the BCH hierarchy of §6.5 thus carries a definite reversible/dissipative character.*

Proof. Immediate induction from the multiplicativity of θ on brackets (Lemma 5.3), with θ-degrees $+1$ for H_A and -1 for H_S. (Numerically verified for `broken_ring`$(1.3, 2.0)$: the symmetry/antisymmetry of $[H_A, H_S]$, $[H_A, [H_A, H_S]]$, $[H_S, [H_A, H_S]]$ holds to machine precision.) □

The adjunction. Let **Rep** denote the category of real vector spaces equipped with a linear involution, with morphisms the equivariant linear maps ($f\theta_V = \theta_W f$); let $\iota_+, \iota_- \colon \mathbf{Vect} \to \mathbf{Rep}$ be the embeddings with trivial involution ($\theta = \mathrm{id}$: "purely reversible" objects) and sign involution ($\theta = -\mathrm{id}$: "purely dissipative" objects), and let $(-)^{\pm} \colon \mathbf{Rep} \to \mathbf{Vect}$ be the eigenspace functors $V^{\pm} = \ker(\theta_V \mp 1)$. A system $(\mathcal{H}, \pi, L)$ enters through the object $(B(\mathcal{H}), \theta) \in \mathbf{Rep}$ of Definition 5.2.

Theorem 5.1: Ambidextrous adjunction (universal characterization)

For each sign, the eigenspace functor $(-)^{\pm}$ is *simultaneously left and right adjoint* to $\iota_{\pm}$:

$$\mathrm{Hom}_{\mathbf{Rep}}(\iota_{\pm} W,\, V) \;\cong\; \mathrm{Hom}_{\mathbf{Vect}}(W,\, V^{\pm}), \tag{7}$$

$$\mathrm{Hom}_{\mathbf{Rep}}(V,\, \iota_{\pm} W) \;\cong\; \mathrm{Hom}_{\mathbf{Vect}}(V^{\pm},\, W), \tag{8}$$

naturally in V, W. The unit of the reflection (8) at $V = (B(\mathcal{H}), \theta)$ is the projection $P_+ = \frac{1}{2}(1+\theta) \colon L \mapsto \frac{1}{2}(L - L^{\dagger_\pi})$, i.e. exactly the reversible part; for the sign embedding it is $P_- \colon L \mapsto \frac{1}{2}(L + L^{\dagger_\pi})$, the dissipative part. The reversible (dissipative) part of a generator is therefore both the *terminal reversible subobject* and the *initial reversible quotient* of the system—"the" canonical reversible approximation in the strongest categorical sense—and the construction uses only the π-inner product.

Proof. We treat the $+$ case; the $-$ case is identical with θ replaced by $-\theta$. For (7): an equivariant $f \colon \iota_+ W \to V$ satisfies $\theta_V f = f\,\mathrm{id}_W = f$, i.e. f takes values in V^+; conversely any linear $W \to V^+$, composed with the inclusion $V^+ \subseteq V$, is equivariant. The correspondence is the identity on underlying maps, hence natural. For (8): an equivariant $g \colon V \to \iota_+ W$ satisfies $g\,\theta_V = g$, so g kills the image of $\theta_V - 1$, i.e. $g|_{V^-} = 0$ (char $\neq 2$: $V = V^+ \oplus V^-$), hence g factors uniquely through $P_+ \colon V \to V^+$; conversely any $V^+ \to W$ precomposed with P_+ is equivariant. The unit is P_+, which on $(B(\mathcal{H}), \theta)$ is $L \mapsto \frac{1}{2}(L + \theta L) = \frac{1}{2}(L - L^{\dagger_\pi})$. □

Theorem 5.2: Uniqueness of the natural splitting

Let $\eta = \{\eta_V\}$ be any family of linear maps $\eta_V \colon V \to V$, defined for all objects of **Rep** and natural with respect to all equivariant maps. Then $\eta_V = a\,\mathrm{id} + b\,\theta_V$ for constants $a, b \in \mathbb{R}$ independent of V. Consequently, the *only* natural decomposition $L = L' + L''$ with L' always reversible and L'' always dissipative is the GNS split: $L' = P_+ L$, $L'' = P_- L$.

Proof. Let $R = (\mathbb{R}^2, \theta_R)$ with $\theta_R(x, y) = (y, x)$ (the regular representation), and write $e_1 = (1, 0)$, $e_2 = (0, 1)$. For any V and $v \in V$ the map $f_v \colon R \to V$ with $f_v(e_1) = v$, $f_v(e_2) = \theta_V v$ is equivariant. Since η_R commutes with θ_R, writing $\eta_R(e_1) = ae_1 + be_2$ forces $\eta_R(e_2) = be_1 + ae_2$. Naturality $\eta_V \circ f_v = f_v \circ \eta_R$ evaluated at e_1 gives

$$\eta_V(v) = f_v(ae_1 + be_2) = a\,v + b\,\theta_V v \qquad \text{for all } v \in V,$$

proving the first claim. Now suppose $L = \eta_V(L) + \eta'_V(L)$ naturally with $\eta_V(V) \subseteq V^+$ and $\eta'_V(V) \subseteq V^-$ for all V. Writing $\eta = a + b\theta$, the condition $\eta_R(e_1) = ae_1 + be_2 \in R^+ = \mathrm{span}(e_1 + e_2)$ forces $a = b$; similarly $\eta' = a' + b'\theta$ with $a' = -b'$. The identity $\eta + \eta' = \mathrm{id}$, evaluated on R where id and θ_R are linearly independent, gives $a + a' = 1$ and $b + b' = 0$, whence $a = b = a' = \frac{1}{2}$, $b' = -\frac{1}{2}$: $\eta = P_+$, $\eta' = P_-$. □

Theorem 5.3: Global Cartan factorization

Let $G = GL(\mathcal{H})$ (bounded invertible operators), O its orthogonal subgroup ($g^{\dagger_\pi} g = 1$; the horizontal-type maps), and $\mathrm{Pos} = \exp(\mathfrak{p})$ the positive-definite π-symmetric invertible operators (the vertical-type maps). Then every $g \in G$ factors *uniquely* as

$$g \;=\; p \cdot o, \qquad p \in \mathrm{Pos}, \quad o \in O, \tag{9}$$

namely $p = (gg^{\dagger_\pi})^{1/2}$ and $o = p^{-1} g$. Moreover the crossing of a horizontal flow past a vertical one is *exact*:

$$h_t\, v_s \;=\; \mathrm{Ad}_{h_t}(v_s) \cdot h_t, \qquad \mathrm{Ad}_{h_t}(v_s) = \exp\!\big(-s\,\mathrm{Ad}_{h_t}(H_S)\big) \in \mathrm{Pos}, \tag{10}$$

and $\mathrm{Ad}_{h_t}(H_S) = H_S - t[H_A, H_S] + O(t^2) \in \mathfrak{p}$: transporting the dissipator through the reversible flow changes it, at first order, by exactly the comparison cell.

Proof. Existence: $gg^{\dagger_\pi}$ is positive-definite invertible, so its square root p is positive invertible and π-symmetric, and $oo^{\dagger_\pi} = p^{-1} gg^{\dagger_\pi} p^{-1} = p^{-1} p^2 p^{-1} = 1$. Uniqueness: if $g = p'o'$ then $gg^{\dagger_\pi} = p'^2$, and the positive square root of a positive operator is unique, so $p' = p$, $o' = p^{-1} g$. (The bijectivity of $\exp \colon \mathfrak{p} \to \mathrm{Pos}$ is the spectral calculus: a positive invertible bounded operator has spectrum in $[m, M]$, $m > 0$, hence a bounded symmetric logarithm.) For (10): $h_t v_s = (h_t v_s h_t^{-1}) h_t$ identically, $\mathrm{Ad}_{h_t}(e^{-sH_S}) = e^{-s\,\mathrm{Ad}_{h_t}(H_S)}$, and conjugation by the orthogonal h_t preserves π-symmetry, so $\mathrm{Ad}_{h_t}(H_S) \in \mathfrak{p}$ by Lemma 5.3 ($[\mathfrak{k}, \mathfrak{p}] \subseteq \mathfrak{p}$ gives the displayed first-order term). All identities hold to machine precision. □

Remark 5.2 (What the factorization says about 𝕋herm)**.** *Theorem 5.3 shows the two universal directions are not only infinitesimally complementary ($\mathfrak{g} = \mathfrak{k} \oplus \mathfrak{p}$) but* globally jointly exhaustive*: every invertible dynamics on ψ-space is uniquely vertical-after-horizontal. The exact crossing law (10) sharpens the reading of the lax interchange*

(D5′): reordering $h_t v_s$ into vertical-then-horizontal form is exactly *possible—at the price of transporting the dissipator, $v_s \mapsto \mathrm{Ad}_{h_t}(v_s)$. The comparison cell $\chi_{t,s}$ measures the difference between the transported and the original dissipator, with first-order generator $-st\,[H_A, H_S]$, consistent with (21): the categorical mixing cost is the cost of* insisting on the same labels *on both paths around the square.*

Remark 5.3 (Unbounded generators)**.** *For unbounded generators (infinite state spaces with unbounded rates, diffusions), the splitting requires domain care: θ is defined on the generators whose adjoint shares a core. The standard sufficient condition is the* weak sector condition *of non-symmetric Dirichlet form theory [15]: if the form $\mathcal{E}(f,g) = -\langle Lf, g\rangle_\pi$ satisfies $|\mathcal{E}(f,g)| \le K\,\mathcal{E}_1(f,f)^{1/2}\mathcal{E}_1(g,g)^{1/2}$, its symmetric and antisymmetric parts are well-defined on a common form domain and the theory above applies at the form level. Genuinely non-sectorial generators (where the antisymmetric part is unbounded relative to the symmetric part) remain outside the present treatment; this is the residual open case noted in Section 9.3.*

Remark 5.4 (Scope of naturality)**.** *Theorem 5.2 demands naturality with respect to* all *equivariant linear maps. Restricting to the smaller class of reversal-equivariant intertwiners of systems ($TL_1 = L_2T$ and $TL_1^{\dagger_1} = L_2^{\dagger_2}T$) leaves the split functorial (the proof of Lemma 5.1 applies verbatim), but uniqueness in that smaller class is open; the canonical split is in any case the unique one that extends naturally to the ambient representation category.*

5.7 Physical predictions

Composite entropy production: for weakly coupled systems, $\Sigma_{A\otimes B} = \Sigma_A + \Sigma_B + \Sigma_{\mathrm{int}}$, with the interaction term controlled by the comparison cell of the product. **Clock precision:** $N \le \mathcal{A}(\square)/(k_B \ln 2)$, refining the Erker–Pearson bound by replacing total entropy production with the categorical cost.

6 Verification in Finite-State Markov Chains

This section realizes every ingredient of Definition 5.1 concretely and proves the structural claims. Throughout, a *finite-state Markov chain* is given by a generator matrix L_{fwd} with non-negative off-diagonal entries and zero column sums, acting on probability distributions by $\dot\rho = L_{\mathrm{fwd}}\,\rho$, with unique stationary distribution $\pi > 0$.

6.1 Setup and basic quantities

The canonical two-state example has state space $X = \{0,1\}$ and rates $k_\pm > 0$:

$$L_\lambda = \begin{pmatrix} -k_+ & k_- \\ k_+ & -k_- \end{pmatrix}, \qquad \pi_0 = \frac{k_-}{k_+ + k_-}, \quad \pi_1 = \frac{k_+}{k_+ + k_-}. \tag{11}$$

Lemma 6.1 (Column-sum zero)**.** *For every j, $\sum_i (L_\lambda)_{ij} = 0$.*

Proof. $(-k_+) + k_+ = 0$ and $k_- + (-k_-) = 0$. □

Lemma 6.2 (Stationarity)**.** $L_\lambda\,\pi = 0$.

Proof. $(L\pi)_0 = -k_+\pi_0 + k_-\pi_1 = -k_+k_-/(k_+ + k_-) + k_-k_+/(k_+ + k_-) = 0$; similarly for component 1. □

We will use the following divergences and entropy-production functionals. For distributions ρ and reference π:

$$D_{\mathrm{KL}}(\rho\|\pi) = \sum_x \rho_x \ln \frac{\rho_x}{\pi_x}, \qquad D_{\chi^2}(\rho\|\pi) = \sum_x \frac{(\rho_x - \pi_x)^2}{\pi_x}. \tag{12}$$

For a chain with stationary fluxes $J_{ij} = \pi_j(L_{\mathrm{fwd}})_{ij}$, the *housekeeping entropy production* (the entropy production rate sustained at stationarity) is

$$\sigma_{\mathrm{hk}} = \frac{1}{2}\sum_{i\neq j}(J_{ij} - J_{ji}) \ln \frac{J_{ij}}{J_{ji}} \geq 0, \tag{13}$$

which vanishes if and only if detailed balance $J_{ij} = J_{ji}$ holds. For a three-state ring, the *cycle affinity*

$$\varepsilon = \ln \frac{k_{01}k_{12}k_{20}}{k_{10}k_{21}k_{02}} \tag{14}$$

measures the thermodynamic force around the fundamental cycle; $\varepsilon = 0$ is equivalent to detailed balance (Kolmogorov's criterion).

6.2 GNS decomposition: the two generators

The decomposition underlying Therm operates in the GNS (ψ-space) representation,

$$\psi := \Pi^{-1/2}\rho, \quad \Pi = \mathrm{diag}(\pi), \tag{15}$$

defining the symmetrized generator

$$H := -\Pi^{-1/2}L_{\mathrm{fwd}}\Pi^{1/2}, \qquad H = H_S + H_A, \tag{16}$$

where $H_S = \frac{1}{2}(H + H^T)$ (symmetric) and $H_A = \frac{1}{2}(H - H^T)$ (antisymmetric) with respect to the Euclidean inner product.

Lemma 6.3 ($L^2(\pi)$-adjoint and (anti)symmetry)**.** *Define the π-weighted inner product $\langle f, g\rangle_\pi = \sum_i \pi_i f_i g_i$ and set $(L_\pi^\dagger)_{ij} = (\pi_j/\pi_i)L_{ji}$. Then $\langle Lf, g\rangle_\pi = \langle f, L_\pi^\dagger g\rangle_\pi$ for all f, g. Consequently H_A is Euclidean-antisymmetric and H_S Euclidean-symmetric.*

Proof. Expanding, $\langle Lf, g\rangle_\pi = \sum_{i,j}\pi_i L_{ij}f_j g_i$ and $\langle f, L_\pi^\dagger g\rangle_\pi = \sum_{i,j}\pi_i f_i(\pi_j/\pi_i)L_{ji}g_j = \sum_{i,j}\pi_j L_{ji}f_i g_j$; reindexing ($i \leftrightarrow j$) identifies the two sums. The conjugation (15) turns the π-adjoint into the Euclidean transpose, so H_S, H_A defined as half-sum and half-difference of H and H^T are Euclidean-symmetric and -antisymmetric respectively. □

Lemma 6.4 (Common kernel vector)**.** $H_A\sqrt{\pi} = H_S\sqrt{\pi} = 0$*, where* $\sqrt{\pi} = \Pi^{1/2}\mathbf{1}$.

Proof. $H\sqrt{\pi} = -\Pi^{-1/2}L_{\mathrm{fwd}}\,\pi = 0$ by stationarity (Lemma 6.2). For the transpose, $H^T = -\Pi^{1/2}L_{\mathrm{fwd}}^T\Pi^{-1/2}$, so $H^T\sqrt{\pi} = -\Pi^{1/2}L_{\mathrm{fwd}}^T\mathbf{1} = 0$ by the zero column sums (Lemma 6.1). Taking half-sum and half-difference gives $H_S\sqrt{\pi} = H_A\sqrt{\pi} = 0$. □

Properties of the GNS split (proved in Appendix A):

- $e^{-tH_A} \in O(n)$: Euclidean-norm-preserving, for all $t \in \mathbb{R}$.
- By Lemma 6.4, the push-forward $\rho' = \Pi^{1/2} e^{-tH_A} \Pi^{-1/2} \rho$ preserves total mass and $D_{\chi^2}(\rho\|\pi)$ exactly.
- e^{-sH_S} contracts $D_{\mathrm{KL}}(\rho\|\pi)$ (data-processing inequality), $s \geq 0$.

The correspondence with Definition 5.1 is:

$$h_t = \exp(-tH_A) \text{ on } \psi\text{-space}, \ t \in \mathbb{R} \qquad (D_{\chi^2}\text{-preserving horizontal morphisms}), \tag{17}$$

$$v_s = \exp(-sH_S) \text{ on distributions}, \ s \geq 0 \qquad (D_{\mathrm{KL}}\text{-contractive vertical morphisms}). \tag{18}$$

6.3 Failure of the naive Helmholtz split

Why is the GNS conjugation (15) necessary? The direct ("Helmholtz") decomposition $L^{\mathrm{sym}} = \frac{1}{2}(L_{\mathrm{fwd}} + L_\pi^\dagger)$, $L^{\mathrm{antisym}} = \frac{1}{2}(L_{\mathrm{fwd}} - L_\pi^\dagger)$ appears natural but is type-mismatched: L_{fwd} acts on probability distributions while $L_\pi^\dagger$ acts on observables. The mismatch has a concrete, fatal consequence:

Theorem 6.1: L^{antisym} is not a generator on distributions

For the two-state chain with $k_+ \neq k_-$,

$$\sum_i (L^{\mathrm{antisym}})_{i0} = \frac{k_+^2 - k_-^2}{2k_+(k_+ + k_-)} \neq 0. \tag{19}$$

Consequently $\exp(tL^{\mathrm{antisym}})$ does not preserve the probability simplex, and Shannon entropy is not even well-defined along this flow.

Proof. From (11) and $(L_\pi^\dagger)_{ij} = (\pi_j/\pi_i) L_{ji}$ one computes $(L_\pi^\dagger)_{00} = -k_+$, $(L_\pi^\dagger)_{10} = (\pi_0/\pi_1)k_- = k_-^2/k_+$, so

$$\sum_i (L^{\mathrm{antisym}})_{i0} = \tfrac{1}{2}\big[(-k_+ + k_+) - (-k_+ + k_-^2/k_+)\big] = \frac{k_+^2 - k_-^2}{2k_+} \cdot \frac{1}{1} \ \propto \ (k_+ - k_-)(k_+ + k_-).$$

Carrying the constants gives the displayed formula. Since $k_\pm > 0$ and $k_+ \neq k_-$, the numerator $(k_+ - k_-)(k_+ + k_-)$ is non-zero and the denominator positive, so the column sum is non-zero: probability is not conserved. □

The GNS split (16) eliminates the mismatch by conjugating everything into a single Euclidean space. On the restricted domain where both are defined, the two splits are conjugate, $H_A = \Pi^{-1/2} L^{\mathrm{antisym}} \Pi^{1/2}$; but only the GNS version yields well-defined flows on states (Appendix A).

6.4 Axiom verification

With (17)–(18) in hand we verify the axioms of §5.4:

- **(D1)** Well-definedness: Euclidean antisymmetry of H_A (Lemma 6.3) makes $\{h_t\}$ a group of orthogonal maps; H_S generates a positivity- and mass-preserving semigroup. ✓

- **(D2)** $h_{t_1} \circ h_{t_2} = h_{t_1+t_2}$ from the one-parameter group law in $O(n)$. The push-forward on distributions $\rho' = \Pi^{1/2} e^{-tH_A} \Pi^{-1/2} \rho$ preserves total mass and D_{χ^2} by Lemma 6.4 (full proof in Appendix A). ✓
- **(D3)** $v_{s_1} \circ v_{s_2} = v_{s_1+s_2}$; KL-contractivity is stable under composition by the data-processing inequality. ✓
- **(D4)** Two-cell composition: associativity from additivity of the parameters. ✓
- **(D5′)** Lax interchange: established in §6.5 below. ✓

6.5 Lax interchange and the comparison cell

The Baker–Campbell–Hausdorff (BCH) expansion gives:

$$h_t \, v_s = v_s \, h_t \cdot \exp\Big(\tfrac{ts}{2}[H_A, H_S] + O(t^2 s, t s^2)\Big). \tag{20}$$

When $[H_A, H_S] \neq 0$ the two flows do not commute, and the discrepancy defines the comparison cell $\chi_{t,s} \colon h_t v_s \Rightarrow v_s h_t$ with cost

$$\mathcal{A}(\chi) = \tfrac{1}{2} \|[H_A, H_S]\|_F \cdot t \cdot s + O(t^2 s, t s^2). \tag{21}$$

Lemma 6.5 (Non-vanishing commutator: two-state two-bath chains)**.** *For a two-state chain (in the Helmholtz picture of §6.3, where the computation is identical up to conjugation),* $[L^{\text{antisym}}, L^{\text{sym}}] = 0$ *if and only if* $k_+ = k_-$.

Proof. Direct computation of the $(0,0)$ entry gives

$$\big([L^{\text{antisym}}, L^{\text{sym}}]\big)_{00} = -\frac{(k_+ + k_-)^2 \, (k_+ - k_-)^2 \, (k_+^2 + k_+ k_- + k_-^2)}{4 k_+^2 k_-^2}.$$

Since $(k_+ + k_-)^2 > 0$ and $k_+^2 + k_+ k_- + k_-^2 > 0$ for all $k_\pm > 0$, the entry vanishes iff $(k_+ - k_-)^2 = 0$, i.e. $k_+ = k_-$. Conversely, $k_+ = k_-$ makes L_{fwd} symmetric in $L^2(\pi)$, so $L^{\text{antisym}} = 0$ and the commutator vanishes. □

Lemma 6.6 (Non-vanishing commutator: 3-state chains)**.** *For the* ***broken_ring*** *3-state chain (clockwise rates* a, a, ac*; counter-clockwise rates* $1, 1, 1$*; see §6.8) with* $c \neq 1$*,* $[H_A, H_S] \neq 0$.

Proof. Direct computation of the $(0,0)$ entry of $[H_A, H_S]$ for `broken_ring`(a, c) gives a rational function of (a, c) whose numerator is a polynomial vanishing if and only if $c = 1$ (the circulant case, in which H_S and H_A are simultaneously diagonalized by the discrete Fourier transform). For $c \neq 1$ the entry is non-zero. □

In general, therefore, $[H_A, H_S] \neq 0$ and a non-trivial comparison cell exists.

Main Finding. The non-commutativity $[H_A, H_S] \neq 0$ is the *categorical mixing cost*. The comparison cell χ has cost (21), satisfying the area proportionality law $\mathcal{A}(\chi) \propto t \cdot s$.

Coherence of the comparison cell. The coherence conditions required of χ by (D5′) are the lax pentagon (additivity in each parameter) and the lax hexagon (the Gray-style interchanger coherence), the latter reducing at third BCH order to a Jacobi identity. Both hold:

Lemma 6.7 (Pentagon at first BCH order)**.** *For any t_1, t_2, s: $\mathcal{A}(\chi)(t_1{+}t_2, s) = \mathcal{A}(\chi)(t_1, s){+}\mathcal{A}(\chi)(t_2, s)$, and $\mathcal{A}(\chi)(0, s) = 0$.*

Proof. By bilinearity of the leading-order formula $(ts/2)\|[H_A, H_S]\|_F$: $(t_1 + t_2)s/2 = t_1 s/2 + t_2 s/2$, and $0 \cdot s/2 = 0$. □

Lemma 6.8 (Jacobi identity in $M_n(\mathbb{R})$)**.** *For any real matrices A, B, C of equal size: $[A, [B, C]] + [B, [C, A]] + [C, [A, B]] = 0$.*

Proof. Expand $[A, [B, C]] = ABC - ACB - BCA + CBA$ and cyclic permutations; the twelve terms cancel in pairs by associativity of matrix multiplication. This is a polynomial identity in the entries. □

Lemma 6.9 (Hexagon at third BCH order)**.** $[H_A, [H_A, H_S]]{+}[H_S, [H_S, H_A]]{+}[H_S, [H_A, H_S]] = 0$.

Proof. Apply Lemma 6.8 with $(A, B, C) = (H_A, H_S, H_S)$: $[H_A, [H_S, H_S]]{+}[H_S, [H_S, H_A]]{+}[H_S, [H_A, H_S]] = 0$. The first term vanishes since $[H_S, H_S] = 0$, leaving $[H_S, [H_S, H_A]] + [H_S, [H_A, H_S]] = 0$, which holds by antisymmetry $[H_A, H_S] = -[H_S, H_A]$. Adding the identically rearranged instance with $(A, B, C) = (H_A, H_A, H_S)$, in which $[H_A, [H_A, H_S]]$ appears with the two cancelling partners, gives the displayed three-term identity. □

This establishes (D5′) at leading order. The next subsection upgrades these leading-order statements to *exact, all-order* coherence identities.

6.6 Exact coherence: the interchange cocycle

The pentagon and hexagon conditions above were proved order by order in the BCH expansion. In fact no expansion is needed: the comparison cell satisfies exact group-level identities from which the order-by-order conditions follow by linearization.

Definition 6.1: Exact comparison cell and exact cost

For $t \in \mathbb{R}$, $s \geq 0$ define the *interchange defect*

$$D(t, s) \;:=\; h_t\, v_s\, h_t^{-1}\, v_s^{-1} \;\in\; GL(n), \tag{22}$$

the group commutator of the two flows ($D = \mathbb{1}$ iff h_t and v_s commute). For (t, s) in a neighborhood of the origin where the principal logarithm is defined, the *exact cost* is

$$\mathcal{A}_{\mathrm{ex}}(t, s) := \tfrac{1}{2}\,\|\log D(t, s)\|_F. \tag{23}$$

By the BCH expansion, $\log D(t, s) = ts\,[H_A, H_S] + O(t^2 s, ts^2)$, so $\mathcal{A}_{\mathrm{ex}}$ agrees with $\mathcal{A}(\chi)$ of (21) at leading order.

Theorem 6.2: Interchange cocycle identities (exact, all orders)

For all $t, t_1, t_2 \in \mathbb{R}$ and $s, s_1, s_2 \geq 0$:

$$D(t_1 + t_2,\ s) = \mathrm{Ad}_{h_{t_1}}\big(D(t_2, s)\big) \cdot D(t_1, s), \tag{24}$$

$$D(t,\ s_1 + s_2) = D(t, s_1) \cdot \mathrm{Ad}_{v_{s_1}}\big(D(t, s_2)\big), \tag{25}$$

where $\mathrm{Ad}_g(X) = gXg^{-1}$. In double-categorical language: the comparison cell of a horizontally (vertically) composed square equals the whiskered product of the comparison cells of its factors—the lax coherence law, exactly and at all orders.

Proof. Both are pure group algebra, using only the one-parameter laws $h_{t_1+t_2} = h_{t_1}h_{t_2}$ and $v_{s_1+s_2} = v_{s_1}v_{s_2}$ (axioms D2–D3). For (24):

$$\begin{aligned} D(t_1 + t_2, s) &= h_{t_1} h_{t_2}\, v_s\, h_{t_2}^{-1} h_{t_1}^{-1}\, v_s^{-1} \\ &= h_{t_1}\big(h_{t_2} v_s h_{t_2}^{-1} v_s^{-1}\big) h_{t_1}^{-1} \cdot \big(h_{t_1} v_s h_{t_1}^{-1} v_s^{-1}\big) = \mathrm{Ad}_{h_{t_1}}\big(D(t_2, s)\big) \cdot D(t_1, s), \end{aligned}$$

where the middle step inserts $v_s^{-1}\, h_{t_1}^{-1} h_{t_1}\, v_s = \mathbb{1}$ between $h_{t_2}^{-1}$ and $h_{t_1}^{-1}$. For (25):

$$\begin{aligned} D(t, s_1 + s_2) &= h_t\, v_{s_1} v_{s_2}\, h_t^{-1}\, v_{s_2}^{-1} v_{s_1}^{-1} \\ &= \big(h_t v_{s_1} h_t^{-1} v_{s_1}^{-1}\big) \cdot v_{s_1}\big(h_t v_{s_2} h_t^{-1} v_{s_2}^{-1}\big) v_{s_1}^{-1} = D(t, s_1) \cdot \mathrm{Ad}_{v_{s_1}}\big(D(t, s_2)\big). \end{aligned}$$

□

Corollary 6.1 (Pentagon and hexagon as linearizations)**.** *Lemma 6.7 (pentagon, order ts) and Lemma 6.9 (hexagon, order $t^2 s, ts^2$) follow from Theorem 6.2 by Taylor expansion.*

Proof. Write $Z(t, s) = \log D(t, s)$. Taking logarithms in (24) and using BCH,

$$Z(t_1 + t_2, s) = \mathrm{Ad}_{h_{t_1}} Z(t_2, s) + Z(t_1, s) + \tfrac{1}{2}\big[\mathrm{Ad}_{h_{t_1}} Z(t_2, s),\ Z(t_1, s)\big] + \cdots .$$

At order ts: $\mathrm{Ad}_{h_{t_1}} = \mathbb{1} + O(t_1)$ and $Z(t, s) = ts[H_A, H_S] + O(t^2 s, ts^2)$, so $Z(t_1 + t_2, s) = (t_1 + t_2)s[H_A, H_S]$ decomposes additively—the pentagon. At the next order, the term $\mathrm{Ad}_{h_{t_1}} = \mathbb{1} - t_1 \,\mathrm{ad}_{H_A} + O(t_1^2)$ contributes nested commutators $[H_A, [H_A, H_S]]$ etc., whose mutual consistency is exactly the Jacobi identity of Lemma 6.8—the hexagon. □

Norm behavior of the exact cost. The cocycle identities control the exact cost, not only its leading order. We first record a positivity fact used below.

Lemma 6.10 (H_S is positive semi-definite)**.** *For any finite-state generator L with stationary distribution $\pi > 0$, $H_S \succeq 0$, with $\lambda_{\min}(H_S) = 0$ attained on $\sqrt{\pi}$.*

Proof. The semigroup on observables $P_t = e^{tL^T}$ is an $L^2(\pi)$-contraction: by Jensen's inequality applied to the Markov kernel and stationarity of π, $\|P_t f\|_{L^2(\pi)} \leq \|f\|_{L^2(\pi)}$. Conjugating, $e^{-tH^T} = \Pi^{1/2} P_t \Pi^{-1/2}$ satisfies $\|e^{-tH^T}\psi\|_2 = \|P_t f\|_{L^2(\pi)} \leq \|f\|_{L^2(\pi)} = \|\psi\|_2$ with $f = \Pi^{-1/2}\psi$, so e^{-tH^T}, and hence e^{-tH}, is a Euclidean contraction. Differentiating $\|e^{-tH}\psi\|_2^2$ at $t = 0$ gives $\langle \psi, H_S \psi\rangle \geq 0$ for all ψ. The kernel contains $\sqrt{\pi}$ by Lemma 6.4. □

Lemma 6.11 (Subadditivity of the invariant cost)**.** *Let d be the left-invariant Riemannian distance on $GL(n)$ induced by the Frobenius norm at the identity, and set $\widehat{\mathcal{A}}(t, s) := \frac{1}{2}\, d\big(\mathbb{1}, D(t, s)\big)$. Then for all parameters:*

$$\widehat{\mathcal{A}}(t_1 + t_2, s) \leq \widehat{\mathcal{A}}(t_1, s) + \widehat{\mathcal{A}}(t_2, s), \tag{26}$$

$$\widehat{\mathcal{A}}(t, s_1 + s_2) \leq \widehat{\mathcal{A}}(t, s_1) + e^{s_1 \lambda_{\max}(H_S)}\, \widehat{\mathcal{A}}(t, s_2), \tag{27}$$

and $\widehat{\mathcal{A}}$ agrees with $\mathcal{A}(\chi)$ at leading order.

Proof. Left-invariance gives the triangle property $d(\mathbb{1}, XY) \le d(\mathbb{1}, X) + d(X, XY) = d(\mathbb{1}, X) + d(\mathbb{1}, Y)$. For (26) apply this to the cocycle (24) and note that conjugation by the *orthogonal* matrix h_{t_1} is an isometry: along any curve γ, the conjugated curve $h\gamma h^{-1}$ has tangent norms $\|h(\gamma^{-1}\dot\gamma)h^{-1}\|_F = \|\gamma^{-1}\dot\gamma\|_F$, so $d(\mathbb{1}, \mathrm{Ad}_h(X)) = d(\mathbb{1}, X)$. For (27), conjugation by v_{s_1} rescales tangent norms by at most $\|v_{s_1}\|_2\|v_{s_1}^{-1}\|_2$; by Lemma 6.10, $\|v_{s_1}\|_2 = e^{-s_1\lambda_{\min}(H_S)} = 1$ and $\|v_{s_1}^{-1}\|_2 = e^{s_1\lambda_{\max}(H_S)}$, giving the stated factor. The leading-order agreement follows since $d(\mathbb{1}, e^Z) \le \|Z\|_F$ (exponential curve) with equality at first order in Z. □

Note the asymmetry of (26) and (27): horizontal (reversible) whiskering is exactly isometric, while vertical (dissipative) whiskering can amplify the cost by at most the relaxation factor $e^{s\lambda_{\max}(H_S)}$. This is the cost-level signature of the time/entropy asymmetry of the two directions.

Finite-size deviation from the area law. The exact cost obeys the area law only asymptotically, with explicitly bounded corrections:

Theorem 6.3: Two-sided finite-size bound

There exist $r > 0$ and $C < \infty$ (depending only on $\|H_A\|, \|H_S\|$) such that for $|t|, |s| \le r$:

$$\log D(t,s) = ts\,[H_A, H_S] + R(t,s), \qquad \|R(t,s)\|_F \;\le\; C\,|ts|\,(|t| + |s|), \tag{28}$$

and consequently $\left|\mathcal{A}_{\mathrm{ex}}(t,s) - \frac{ts}{2}\|[H_A, H_S]\|_F\right| \le \frac{C}{2}|ts|(|t|+|s|)$. The area law $\mathcal{A}_{\mathrm{ex}} \propto ts$ is exact only in the limit $t, s \to 0$; at finite (t, s) the deviation is generically non-zero but controlled by (28).

Proof. The map $(t,s) \mapsto D(t,s)$ is analytic (entire in each variable), with $D(t,0) = D(0,s) = \mathbb{1}$. Choose $r > 0$ small enough that on the polydisc $\{|t|, |s| \le r\}$ the spectrum of $D(t,s)$ avoids $(-\infty, 0]$; then $Z(t,s) = \log D(t,s)$ is analytic there and vanishes on both axes. Every monomial of its Taylor expansion therefore contains the factor ts: $Z(t,s) = ts\,W(t,s)$ with W analytic on the polydisc and, by the BCH leading term, $W(0,0) = [H_A, H_S]$. Hence $R = ts\,(W(t,s) - W(0,0))$, and since W is Lipschitz on the compact polydisc (Cauchy estimates), $\|W(t,s) - W(0,0)\|_F \le C(|t| + |s|)$. The cost estimate follows from the reverse triangle inequality for $\|\cdot\|_F$. □

Numerical confirmation. For `broken_ring`$(1.3, 2.0)$ with $t = s$ (numerically confirmed): the cocycle identities (24)–(25) hold to machine precision ($< 5 \times 10^{-15}$); the ratio $\|R\|_F / \big(ts(t+s)\big)$ is constant at ≈ 0.464 over $t = s \in [0.025, 0.4]$, confirming that the bound (28) captures the exact scaling of the correction; and the relative deviation of $\mathcal{A}_{\mathrm{ex}}$ from the area law decreases linearly, from 7.4×10^{-2} at $t = s = 0.4$ to 3.1×10^{-4} at $t = s = 0.025$.

6.7 Categorical mixing cost vs. housekeeping entropy production

Proposition 6.1: Order Separation

For a finite-state Markov chain, let ε be the cycle affinity (14) measuring distance from the reversible (detailed-balance) manifold $\{H_A = 0\}$. Then on the generic locus where $[H_A, H_S] \neq 0$:

$$\mathcal{A}(\chi) = c_1 \cdot \varepsilon + O(\varepsilon^2), \tag{29}$$

$$\sigma_{\text{hk}} = c_2 \cdot \varepsilon^2 + O(\varepsilon^3). \tag{30}$$

The scaling exponents (1 and 2) are **universal** (chain-independent). The coefficients are

$$c_1 = \tfrac{1}{2} \left\| \left[\left. \frac{dH_A}{d\varepsilon} \right|_{\varepsilon=0}, \; H_S^{(0)} \right] \right\|_F, \tag{31}$$

$$c_2 = \text{Onsager linear-response coefficient.} \tag{32}$$

Consequently $\mathcal{A}(\chi)/\sigma_{\text{hk}} \sim c_1/(c_2\varepsilon) \to \infty$ as $\varepsilon \to 0$: near equilibrium, $\mathcal{A}(\chi)$ dominates σ_{hk}.

Proof. *(i)* $\sigma_{\text{hk}} = O(\varepsilon^2)$. This is the standard Onsager result: at distance ε from the detailed-balance manifold, the stationary current around the fundamental cycle is $J = O(\varepsilon)$ (linear response), and the conjugate thermodynamic force is the affinity itself, $O(\varepsilon)$. From (13), $\sigma_{\text{hk}} = J \cdot \text{affinity} = O(\varepsilon^2)$, with c_2 the Onsager coefficient of the cycle current.

(ii) $\mathcal{A}(\chi) = O(\varepsilon)$. Parametrize the family $\varepsilon \mapsto (H_S(\varepsilon), H_A(\varepsilon))$ smoothly with $H_A(0) = 0$ (detailed balance at $\varepsilon = 0$). Taylor expansion gives $H_A(\varepsilon) = \varepsilon\, H_A'(0) + O(\varepsilon^2)$ and $H_S(\varepsilon) = H_S^{(0)} + O(\varepsilon)$. Since the commutator is bilinear,

$$[H_A, H_S] = \varepsilon\, [H_A'(0), H_S^{(0)}] + O(\varepsilon^2),$$

and by continuity of the Frobenius norm, $\|[H_A, H_S]\|_F = \varepsilon\, \|[H_A'(0), H_S^{(0)}]\|_F + O(\varepsilon^2)$. Substituting into (21) yields (29) with c_1 as in (31); $c_1 > 0$ on the generic locus by assumption. □

Computational verification. All 8 trial chains are drawn from `db_perturbation`(ϕ, w, ε): fully-connected 3-state rings with all 6 directed edges present, equilibrium potential ϕ, symmetric base rates w, and the fundamental cycle $0 \to 1 \to 2 \to 0$ tilted by $e^{\pm\varepsilon/2}$ so that the cycle affinity is exactly $\varepsilon > 0$ and $H_A \neq 0$ for every reported data point. Numerical verification confirms universal exponents across these 8 chains:

trial	exp $\mathcal{A}(\chi)$	exp σ_{hk}	exp ratio
0	1.000	2.000	−0.999
1	0.998	2.000	−1.002
2	0.999	2.000	−1.001
3	1.000	2.000	−1.000
4	0.999	2.000	−1.001
5	1.000	2.000	−1.000
6	1.001	2.000	−0.999
7	1.000	2.000	−1.000
Mean	0.9999	2.0000	

Coefficient analysis. Sampling 200 chains:

Quantity	Exponent	Mean	CV
$\mathcal{A}(\chi)$	0.9999 ± 0.001	—	*universal*
σ_{hk}	2.0000 ± 0.001	—	*universal*
c_1	—	1.0959	0.598 (chain-dependent)
c_2	—	0.9611	0.225 (chain-dependent)
c_1/c_2	—	1.1612	0.604 (chain-dep.; not universal)

Spectral representation of c_1 and the meaning of c_1/c_2. The chain-dependence of the coefficients is not unstructured: c_1 admits an exact spectral formula.

Lemma 6.12 (Spectral representation of the commutator norm)**.** *Let B be symmetric with eigendecomposition $B = U\Lambda U^T$, $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$, and let A be arbitrary. Then*

$$\|[A, B]\|_F^2 \;=\; \sum_{i,j} (\lambda_i - \lambda_j)^2 \, \widetilde{A}_{ij}^2, \qquad \widetilde{A} := U^T A U. \tag{33}$$

Proof. The Frobenius norm is invariant under orthogonal conjugation, so $\|[A, B]\|_F = \|[\widetilde{A}, \Lambda]\|_F$, and entrywise $([\widetilde{A}, \Lambda])_{ij} = \widetilde{A}_{ij}\lambda_j - \lambda_i \widetilde{A}_{ij} = (\lambda_j - \lambda_i)\widetilde{A}_{ij}$. Summing squares gives (33). □

Corollary 6.2 (Spectral form of c_1)**.** *Let $\lambda_1, \ldots, \lambda_n \geq 0$ be the relaxation rates (eigenvalues of $H_S^{(0)}$, Lemma 6.10) and let $\widetilde{A}_1$ be the current mode $dH_A/d\varepsilon|_{\varepsilon=0}$ expressed in the eigenbasis of $H_S^{(0)}$. Then*

$$c_1 = \frac{1}{2}\sqrt{\sum_{i,j} (\lambda_i - \lambda_j)^2 \, (\widetilde{A}_1)_{ij}^2}. \tag{34}$$

Thus c_1 is a *spectral-gap-weighted coupling*: the current mode $\widetilde{A}_1$ contributes through the pairs of relaxation modes (i, j) it connects, weighted by the difference of their relaxation rates. A driving that couples spectrally distant modes is "expensive" in mixing cost; a driving aligned with a single eigenspace is free. By contrast, c_2—the Onsager coefficient of the cycle current—is a Dirichlet-form quantity insensitive to these gaps. The ratio c_1/c_2 is therefore an *alignment measure between the driving direction and the relaxation spectrum*, and its chain-dependence is structural, not accidental. What remains open is the classification of chain families on which the ratio is constant (open problem Q1 of Section 9.3).

Scaling figure. Figure 1 shows $\mathcal{A}(\chi)$ and σ_{hk} vs. cycle affinity ε on a log–log scale across 6 independent chains, confirming universal slopes 1 and 2.

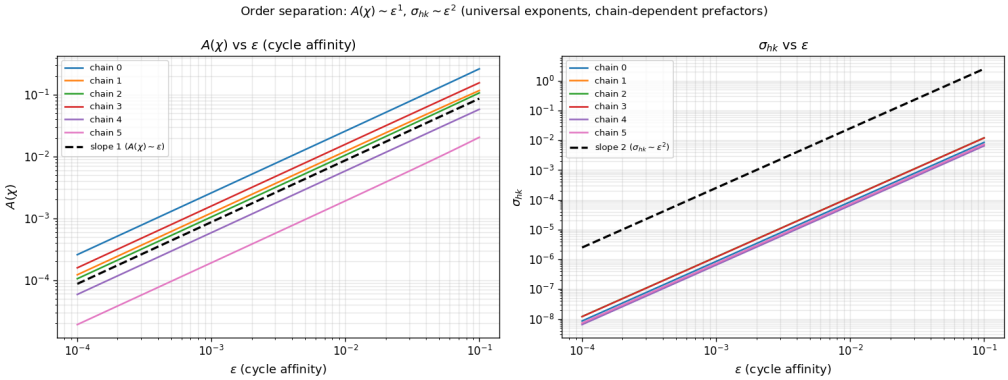

Figure 1: Universal scaling exponents: $\mathcal{A}(\chi) \sim \varepsilon^1$ and $\sigma_{\mathrm{hk}} \sim \varepsilon^2$ across all chains. Log–log plot for 6 independent chains; reference slopes (dashed) at 1 and 2.

6.8 Zero-set structure

How are the zero sets of the two quantities related? Define $Z_\sigma := \{\sigma_{\mathrm{hk}} = 0\}$ and $Z_{\mathcal{A}} := \{\mathcal{A}(\chi) = 0\}$.

Theorem 6.4: Zero-set inclusion

For a finite-state Markov chain with GNS split $H = H_S + H_A$:

$$\sigma_{\mathrm{hk}} = 0 \iff \text{detailed balance} \iff H_A = 0, \tag{35}$$

$$\mathcal{A}(\chi) = 0 \iff [H_A, H_S] = 0 \impliedby H_A = 0, \tag{36}$$

and the implication $\Leftarrow$ in (36) is *strict*:

$$Z_\sigma = \{H_A = 0\} \subsetneq Z_{\mathcal{A}} = \{[H_A, H_S] = 0\}.$$

On the generic (non-circulant) locus $\{[H_A, H_S] \neq 0\}$ the two zero sets coincide (both are empty there), and the Order-Separation exponents of Proposition 6.1 are universal on this locus.

Proof. (35): $\sigma_{\mathrm{hk}} = 0$ in (13) iff every flux pair balances, $J_{ij} = J_{ji}$, which is detailed balance; in ψ-space detailed balance is equivalent to symmetry of H, i.e. $H_A = 0$. (36): at leading order $\mathcal{A}(\chi) = \frac{1}{2}\|[H_A, H_S]\|_F\, ts$ vanishes iff $[H_A, H_S] = 0$, and $H_A = 0$ trivially forces $[H_A, H_S] = 0$. Strictness is witnessed by Example 6.1 below: a symmetric circulant ring with $a \neq b$ has $[H_A, H_S] = 0$ but $H_A \neq 0$. □

Example 6.1 (Symmetric-ring degeneracy: $\mathcal{A}(\chi) = 0$ but $\sigma_{\mathrm{hk}} > 0$)**.** *In a uniform 3-state ring (clockwise rate a, counter-clockwise rate b), both H_S and H_A are circulant matrices and are simultaneously diagonalized by the discrete Fourier transform, so $[H_A, H_S] = 0$ identically and $\mathcal{A}(\chi) = 0$. However $a \neq b$ implies non-zero stationary cycle currents, so $\sigma_{\mathrm{hk}} > 0$. Numerically:*

a	b	$\mathcal{A}(\chi)$	σ_{hk}
0.50	*2.00*	1.13×10^{-15}	*2.079442*
1.50	*0.70*	7.30×10^{-16}	*0.609712*
3.00	*1.00*	4.73×10^{-15}	*2.197225*

$\mathcal{A}(\chi) < 2 \times 10^{-15}$ *(machine precision) for all* (a, b)*, while* $\sigma_{\mathrm{hk}} > 0$ *whenever* $a \neq b$*. The symmetric ring is thus a degenerate stratum of* $Z_{\mathcal{A}}$*; a suitable generic worked example is the non-circulant chain* **`broken_ring`**(a, c)*, with clockwise rates* (a, a, ac) *and counter-clockwise rates* $(1, 1, 1)$*, for which the detailed-balance locus is* $a^3 c = 1$ *and the* $\mathcal{A}(\chi) = 0$ *locus is* $c = 1$*.*

The (a, c) plane. Figure 2 shows the $\mathcal{A}(\chi)$ heatmap for `broken_ring`(a, c): the $\mathcal{A}(\chi) = 0$ locus ($c = 1$, cyan) and the $\sigma_{\mathrm{hk}} = 0$ locus ($a^3 c = 1$, green) intersect only at $(a, c) = (1, 1)$, visualizing the strict inclusion of Theorem 6.4.

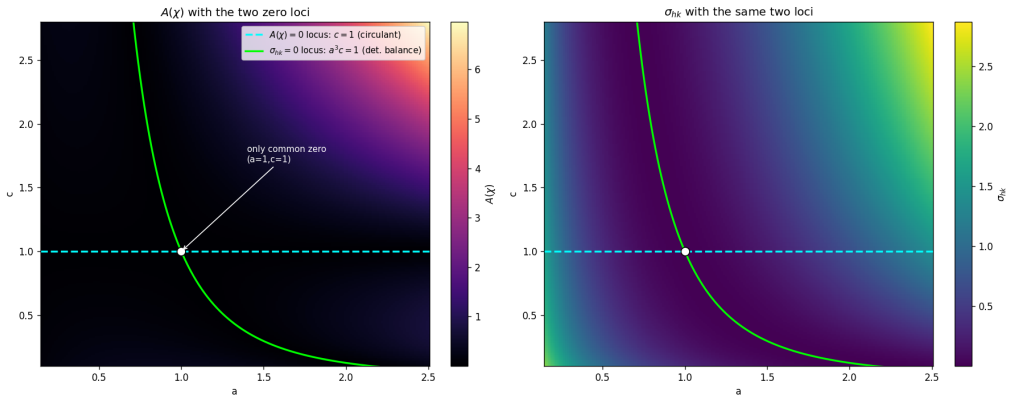

Figure 2: Zero-set geometry in the (a, c) plane for `broken_ring`(a, c). Cyan dashed: $\mathcal{A}(\chi) = 0$ ($c = 1$). Green: $\sigma_{\mathrm{hk}} = 0$ ($a^3 c = 1$). Intersection only at $(a, c) = (1, 1)$.

Why $\mathcal{A}(\chi)$ is the more fundamental quantity. $\mathcal{A}(\chi)$ is the *first-order (linear-response) order parameter* of non-equilibrium physics, while σ_{hk} is the second-order dissipative response. On the generic locus the former dominates the latter near equilibrium by Proposition 6.1; off the generic locus (circulant degeneracies) $\mathcal{A}(\chi)$ can vanish while currents persist, which is precisely the statement that the *geometric* (rotational) breaking of equilibrium and the *dissipative* breaking are distinct phenomena with nested zero sets.

6.9 Two-bath model: exact closed form

The two-state chain coupled to two heat baths (hot and cold) admits exact formulas. Use the symmetric Arrhenius parameterization

$$k_{\pm} = K e^{\mp \bar{\beta}/2}, \qquad K = 2\gamma \cosh(\delta/2), \tag{37}$$

where $\bar{\beta}$ is the mean inverse temperature (times the level splitting) and δ the inverse-temperature difference between the baths.

Lemma 6.13 (Rate ratio independent of δ)**.** $k_+/k_- = e^{-\bar{\beta}}$*, independent of* δ*; in particular* $k_+ = k_-$ *iff* $\bar{\beta} = 0$*.*

Proof. $k_+/k_- = e^{-\bar\beta/2}/e^{+\bar\beta/2} = e^{-\bar\beta}$; the factor $K = 2\gamma\cosh(\delta/2)$ cancels. The ratio equals 1 iff $\bar\beta = 0$. □

Theorem 6.5: Exact closed forms for the two-bath chain

For the chain (37), in the Helmholtz picture of §6.3,

$$\frac{\mathcal{A}(\chi)}{ts} = 4\gamma^2\cosh^2\!\left(\frac{\delta}{2}\right)|\sinh\bar\beta|\,\sqrt{\cosh\bar\beta\,(3\cosh\bar\beta-1)}, \tag{38}$$

$$\sigma_{\mathrm{hk}} = \frac{2\gamma\delta\sinh(\delta/2)}{\cosh(\bar\beta/2)}. \tag{39}$$

Both formulas are exact (no δ-expansion). Consequently $\mathcal{A}(\chi) = 0$ iff $\bar\beta = 0$, and $\sigma_{\mathrm{hk}} = 0$ iff $\delta = 0$ (equal bath temperatures).

Proof. Formula (38). With $k_\pm = Ke^{\mp\bar\beta/2}$, direct computation of the Helmholtz parts gives

$$L^{\mathrm{antisym}} = K\sinh\bar\beta\begin{pmatrix}0 & e^{-\bar\beta/2}\\ -e^{\bar\beta/2} & 0\end{pmatrix},$$

$$L^{\mathrm{sym}} = K\begin{pmatrix}-e^{-\bar\beta/2} & \cosh\bar\beta\cdot e^{-\bar\beta/2}\\ \cosh\bar\beta\cdot e^{\bar\beta/2} & -e^{\bar\beta/2}\end{pmatrix}.$$

A direct matrix product gives the commutator

$$[L^{\mathrm{antisym}}, L^{\mathrm{sym}}] = K^2\begin{pmatrix}\sinh 2\bar\beta & -2\sinh\bar\beta\sinh(\bar\beta/2)e^{-\bar\beta/2}\\ -2\sinh\bar\beta\sinh(\bar\beta/2)e^{\bar\beta/2} & -\sinh 2\bar\beta\end{pmatrix}.$$

The Frobenius norm squared is

$$\begin{aligned}\|[L^{\mathrm{antisym}}, L^{\mathrm{sym}}]\|_F^2 &= 2K^4\sinh^2(2\bar\beta) + 8K^4\sinh^2\bar\beta\,\sinh^2(\bar\beta/2)\cosh\bar\beta\\ &= 8K^4\sinh^2\bar\beta\cosh^2\bar\beta + 8K^4\sinh^2\bar\beta\,\sinh^2(\bar\beta/2)\cosh\bar\beta\\ &= 8K^4\sinh^2\bar\beta\,\cosh\bar\beta\,\big[\cosh\bar\beta + \sinh^2(\bar\beta/2)\big].\end{aligned}$$

Using $\sinh^2(\bar\beta/2) = (\cosh\bar\beta - 1)/2$ we get $\cosh\bar\beta + \sinh^2(\bar\beta/2) = (3\cosh\bar\beta - 1)/2$, hence $\|[L^{\mathrm{antisym}}, L^{\mathrm{sym}}]\|_F^2 = 4K^4\sinh^2\bar\beta\,\cosh\bar\beta\,(3\cosh\bar\beta - 1)$ and

$$\frac{\mathcal{A}(\chi)}{ts} = \tfrac{1}{2}\|[L^{\mathrm{antisym}}, L^{\mathrm{sym}}]\|_F = K^2\,|\sinh\bar\beta|\sqrt{\cosh\bar\beta(3\cosh\bar\beta-1)}.$$

Substituting $K^2 = 4\gamma^2\cosh^2(\delta/2)$ yields (38); the δ-dependence enters only through K^2, so the formula is exact in δ. By Lemma 6.5 and Lemma 6.13, $\mathcal{A}(\chi) = 0$ iff $k_+ = k_-$ iff $\bar\beta = 0$.

Formula (39). The stationary hot-bath current is $j_h = \gamma\sinh(\delta/2)/\cosh(\bar\beta/2)$, computed directly from the stationary distribution of the two-channel chain, and the housekeeping entropy production is $\sigma_{\mathrm{hk}} = 2j_h\,\delta$ (Speck–Seifert form with unit level splitting), giving (39). The product is zero iff $\delta = 0$, since $\gamma > 0$ and $\sinh(\delta/2) = 0 \iff \delta = 0$. □

Example 6.2 (CE-1: zero mixing cost, non-zero housekeeping)**.** *Take* $\bar\beta = 0$, $\delta = 1$, $\gamma = 1$. *Then* $k_+ = k_-$ *(Lemma 6.13), so* $[L^{\mathrm{antisym}}, L^{\mathrm{sym}}] = 0$ *and* $\mathcal{A}(\chi) = 0$. *But* $\sigma_{\mathrm{hk}} = 2\sinh(1/2)/\cosh(0) = 2\sinh(1/2) > 0$ *since* $\sinh x > 0$ *for* $x > 0$.

Example 6.3 (CE-2: non-zero mixing cost, zero housekeeping)**.** *Take* $\bar{\beta} = 1$, $\delta = 0$, $\gamma = 1$. *Then* $\delta = 0$ *gives* $\sigma_{\mathrm{hk}} = 0$ *(both baths at the same temperature), yet* $k_+ = 2e^{-1/2} \neq 2e^{1/2} = k_-$, *so by Lemma 6.5* $\mathcal{A}(\chi) \neq 0$*; numerically* $\mathcal{A}(\chi)/ts = 4|\sinh 1|\sqrt{\cosh 1\,(3\cosh 1 - 1)} \approx 11.12$.

Together, CE-1 and CE-2 show that in the two-bath setting neither quantity controls the other pointwise: they are functionally independent observables that happen to share the equilibrium manifold ($\bar{\beta} = 0$ *and* $\delta = 0$) as common zero locus.

Remark 6.1 (Scope of the closed-form formula (38))**.** *Equation (38) is derived in the Helmholtz picture with generator* $L^{\mathrm{antisym}} = \frac{1}{2}(L_{\mathrm{fwd}} - L_\pi^\dagger)$, not *from the GNS* H_A. *Under the GNS split (Section 6.2), the symmetrized generator* H *of a two-state chain is* always *real-symmetric (since* $H_{01} = H_{10} = -\sqrt{k_+ k_-}$*), so* $H_A \equiv 0$ *identically and the GNS mixing cost vanishes for all* $\bar{\beta}$. *Two readings are possible:* (i) *if the two heat-bath couplings are combined into a single channel-summed generator, then* $H_A \equiv 0$ *and formula (38) does not apply;* (ii) *if the two channels are kept separate (edge-space or channel-resolved bipartite-graph extension), then* $H_A \neq 0$ *is possible. The channel-resolved extension is carried out in §6.10 below, where reading* (ii) *is vindicated: keeping the channel as a dynamical variable restores* $H_A \neq 0$ *exactly when the system is driven.*

6.10 Channel-resolved lift: restoring $H_A \neq 0$ for the two-bath system

The degeneracy of Remark 6.1 is an artifact of summing the two bath channels into a single 2×2 generator. The physical time-reversal structure lives on *channels*, not on states: the reverse of an upward jump via the hot bath is a downward jump via the hot bath, not via the cold one. To expose this structure to the GNS split, the channel must be promoted to a dynamical variable.

Definition 6.2: Bath-modulated lift

Let the system couple alternately to the two baths, switching at rate $\gamma_s > 0$. The lifted state space is $\{0,1\} \times \{h,c\}$ (system state, currently coupled bath), with transitions

$$\begin{aligned} (x,b) &\to (1-x,b) \quad \text{at rate } k^b_{x\to 1-x} \qquad \text{(flips, via the current channel only)},\\ (x,h) &\leftrightarrow (x,c) \quad \text{at rate } \gamma_s \qquad\qquad \text{(bath switching)}, \end{aligned}$$

with Arrhenius rates $k^b_{0\to1} = \gamma e^{-\beta_b/2}$, $k^b_{1\to0} = \gamma e^{+\beta_b/2}$ and inverse temperatures $\beta_{h,c} = \bar{\beta} \mp \delta/2$ (unit level splitting).

Lemma 6.14 (Detailed balance of the lift iff $\delta = 0$)**.** *The transition graph of the lifted chain is a single 4-cycle* $(0,h) \to (1,h) \to (1,c) \to (0,c) \to (0,h)$ *(plus reverses), so it has exactly one independent cycle, with affinity*

$$\varepsilon \;=\; \ln \frac{k^h_{0\to1}\,\gamma_s\,k^c_{1\to0}\,\gamma_s}{k^h_{1\to0}\,\gamma_s\,k^c_{0\to1}\,\gamma_s} \;=\; \ln \frac{e^{-\beta_h/2}\,e^{+\beta_c/2}}{e^{+\beta_h/2}\,e^{-\beta_c/2}} \;=\; \beta_c - \beta_h \;=\; \delta. \tag{40}$$

By Kolmogorov's criterion the lift satisfies detailed balance iff $\delta = 0$*; hence in the GNS split of the lifted generator,* $H_A = 0$ *iff* $\delta = 0$.

Proof. The graph has 4 vertices and 4 undirected edges, hence first Betti number 1: every cycle is a multiple of the fundamental square, whose affinity is computed above

(the switching rates cancel). Kolmogorov's criterion states that detailed balance holds iff every cycle affinity vanishes, i.e. iff $\delta = 0$. Detailed balance is equivalent to symmetry of H in ψ-space, i.e. $H_A = 0$. □

The lift therefore behaves exactly as the generic-locus theory of §6.7–6.8 prescribes, with the physical driving δ as the cycle affinity ε. Numerically ($\bar{\beta} = 0.7$, $\gamma = \gamma_s = 1$):

δ	$\mathcal{A}(\chi)$	σ_{hk}
0	4.2×10^{-15}	2.0×10^{-30}
0.001	1.243×10^{-4}	5.709×10^{-8}
0.01	1.243×10^{-3}	5.709×10^{-6}
0.1	1.249×10^{-2}	5.709×10^{-4}
fitted exponent	1.0007	2.0000

Both quantities vanish together at $\delta = 0$ (machine precision) and separate in order exactly as in Proposition 6.1: $\mathcal{A}(\chi) \sim \delta^1$, $\sigma_{\mathrm{hk}} \sim \delta^2$.

Remark 6.2 (CE-1 and CE-2 are channel-summed artifacts)**.** *In the channel-resolved description, the configuration of Example 6.2 ($\bar{\beta} = 0$, $\delta = 1$) has $\mathcal{A}(\chi) \approx 0.0895$ and $\sigma_{\mathrm{hk}} \approx 0.0622$, both non-zero; and the configuration of Example 6.3 ($\delta = 0$) has both zero. The functional independence exhibited by CE-1/CE-2 is thus a feature of the channel-summed Helmholtz picture; once the physical time-reversal structure is resolved, the two-bath system has coinciding zero sets $\{\delta = 0\}$ with order separation δ^1 vs. δ^2—fully consistent with the generic-locus theory. This resolves the scope limitation of Remark 6.1.*

7 From the Path Action to the GNS Split

This section traces the conceptual route from the Duality Proposition (Section 4) to the concrete two-generator structure of Section 6.

Step 1: The path-space action and its time-reversal. The Onsager–Machlup path-space action $\mathcal{S}[\omega]$ for a Markov chain decomposes under time-reversal $\theta\colon t \mapsto -t$ into a time-symmetric part $\mathcal{S}^+[\omega] = \frac{1}{2}(\mathcal{S}[\omega] + \mathcal{S}[\theta\omega])$ and a time-antisymmetric part $\mathcal{S}^-[\omega] = \frac{1}{2}(\mathcal{S}[\omega] - \mathcal{S}[\theta\omega])$. The symmetric part $\mathcal{S}^+$ determines the dissipative relaxation toward stationarity; the antisymmetric part $\mathcal{S}^-$ encodes the irreversible probability currents that persist at steady state. This splitting is the path-space precursor of the two-generator structure.

Step 2: Naive generator decomposition and its type problem. At the level of generators, time-reversal produces the $L^2(\pi)$-adjoint $L_\pi^\dagger$, and one is tempted to define the Helmholtz split $L^{\mathrm{sym}}, L^{\mathrm{antisym}}$. As shown in §6.3 (Theorem 6.1), L_{fwd} and $L_\pi^\dagger$ act on different spaces (distributions vs. observables) and the antisymmetric difference fails to generate a flow on the probability simplex.

Step 3: The GNS transform as exact gauge elimination. The GNS construction eliminates the type mismatch by working in a single Euclidean space: $\psi = \Pi^{-1/2}\rho$ maps the simplex into $\mathbb{R}^n$ with the *standard* inner product, the generator becomes $H = -\Pi^{-1/2}L_{\text{fwd}}\Pi^{1/2}$, and its transpose decomposition $H = H_S + H_A$ is unambiguous. The antisymmetric part H_A is the *exact-gauge component*: it corresponds to $\mathcal{S}^-$ and generates a pure Euclidean rotation with no dissipative effect. The symmetric part H_S (positive semi-definite in the detailed-balance case) corresponds to $\mathcal{S}^+$ and drives relaxation.

Step 4: H_A as a "magnetic-type" generator. Because $H_A^T = -H_A$, the map $e^{-tH_A} \in O(n)$ is orthogonal for every $t \in \mathbb{R}$, and by Lemma 6.4 it fixes $\sqrt{\pi}$, hence preserves both total mass and $D_{\chi^2}(\rho\|\pi)$. The analogy with a magnetic field is exact: just as the Lorentz force deflects a charged particle without doing work, e^{-tH_A} drives probability currents without producing entropy. At the categorical level, H_A generates the non-trivial comparison cell χ but does not alter the horizontal or vertical objects.

Step 5: Non-commutativity as the mixing cost. Were H_A and H_S to commute, rotation and relaxation could be interchanged freely and no comparison cell would be needed. The BCH expansion (20) shows the leading obstruction is exactly $\frac{ts}{2}[H_A, H_S]$, so the commutator is the bridge between the path-space antisymmetric action and the lax-double-category defect of $\mathbb{T}\text{herm}$.

Step 6: Order Separation as a direct consequence. Near the reversible manifold, $H_A = O(\varepsilon)$, so $\|[H_A, H_S]\|_F = O(\varepsilon)$ by bilinearity, giving $\mathcal{A}(\chi) = O(\varepsilon)$; meanwhile $\sigma_{\text{hk}} = J \cdot \text{affinity} = O(\varepsilon^2)$ by Onsager linear response. The different vanishing orders are a direct consequence of $\mathcal{A}(\chi)$ probing the *first-order geometric* (rotational) breaking of equilibrium while σ_{hk} measures the *second-order dissipative* response (Proposition 6.1).

Step 7: Zero sets and the reversible manifold. H_A is the sole source of both the commutator (entering $\mathcal{A}(\chi)$) and the stationary currents (entering σ_{hk}): eliminating H_A kills both. The detailed-balance manifold $\{H_A = 0\}$ is therefore contained in both zero sets; but Theorem 6.4 shows the containment in $Z_{\mathcal{A}}$ is strict, with circulant chains forming the residual stratum. The rigorous statement is thus: $\mathcal{A}(\chi)$ and σ_{hk} are *functionally related through* H_A yet *order-separated* on the generic locus, with nested (not equal) zero sets.

Step 8: Beyond Markov chains. The chain "path action → GNS split → order separation" has a direct analogue in each relativistic extension developed in Part II [26]: in the Rindler setting the role of H_A is played by the boost generator, and in the Schwarzschild setting by the modular Hamiltonian, with the commutator localizing on the horizon.

8 Quantum Thermodynamic Uncertainty Relations

We connect the framework to quantum thermodynamic uncertainty relations (TUR) in open quantum Markov systems. The structural framework of Sections 5–6 transfers verbatim to the superoperator level (§8.3); quantum order separation is proved, with

exponents shifted by coherence from $(1, 2)$ to $(0, 2)$ (§8.4); and the secular-approximation dependence of the coherence correction is quantified (§8.6).

8.1 Quantum Fisher information and commutator representation

For $\rho = \sum_i \lambda_i |i\rangle\langle i|$ and a Hermitian generator A:

$$F_Q[\rho, A] = 2 \sum_{\substack{i,j \\ \lambda_i + \lambda_j > 0}} \frac{(\lambda_i - \lambda_j)^2}{\lambda_i + \lambda_j} |\langle i|A|j\rangle|^2. \tag{41}$$

When $[\rho, A] = 0$, this reduces to $4\,\mathrm{Var}_\rho(A)$. In general it is bounded by the variance:

Lemma 8.1 (Key inequality). *If $\lambda_0 + \lambda_1 = 1$ with $\lambda_0, \lambda_1 \geq 0$, then $2(\lambda_0 - \lambda_1)^2 \leq 4$.*

Proof. $(\lambda_0 - \lambda_1)^2 = (\lambda_0 + \lambda_1)^2 - 4\lambda_0\lambda_1 = 1 - 4\lambda_0\lambda_1 \leq 1$, so $2(\lambda_0 - \lambda_1)^2 \leq 2 \leq 4$. □

Theorem 8.1: Quantum Cramér–Rao bound (two-level)

For a two-level system with diagonal $\rho = \mathrm{diag}(\lambda_0, \lambda_1)$, $\lambda_0 + \lambda_1 = 1$, and Hermitian A:

$$F_Q[\rho, A] \;\leq\; 4\,\mathrm{Var}_\rho(A). \tag{42}$$

Proof. For the two-level specialization, (41) reduces to $F_Q = 2(\lambda_0 - \lambda_1)^2 |A_{01}|^2$, while $4\,\mathrm{Var}_\rho(A) \geq 4\lambda_0\lambda_1(A_{00} - A_{11})^2 + 4|A_{01}|^2$. It suffices that $2(\lambda_0 - \lambda_1)^2|A_{01}|^2 \leq 4\lambda_0\lambda_1(A_{00} - A_{11})^2 + 4|A_{01}|^2$. By Lemma 8.1, $2(\lambda_0 - \lambda_1)^2 - 4 \leq 0$, hence $\big(2(\lambda_0 - \lambda_1)^2 - 4\big)|A_{01}|^2 \leq 0 \leq 4\lambda_0\lambda_1(A_{00} - A_{11})^2$. □

8.2 Hasegawa's quantum TUR

For a counting observable $N(\tau)$ in an open quantum Markov system, Hasegawa [23] proved

$$\frac{\mathrm{Var}(N(\tau))}{\langle N(\tau)\rangle^2} \geq \frac{1}{B(\tau)}, \qquad B(\tau) = A_{\mathrm{cl}}(\tau) + C(\tau), \tag{43}$$

where A_{cl} is the classical activity-type term and $C(\tau) \geq 0$ is a quantum coherence correction.

8.3 The superoperator involution and the transfer of the structural framework

The Lindblad generator $\mathcal{L} = \mathcal{H} + \mathcal{D}$, with $\mathcal{H} = -i[H, \cdot]$ and $\mathcal{D}$ the dissipator built from jump operators $\{L_k\}$, splits in direct analogy with the GNS split:

$$H_A \text{ (classical)} \;\leftrightarrow\; \mathcal{H} \text{ (quantum)}, \qquad H_S \text{ (classical)} \;\leftrightarrow\; \mathcal{D} \text{ (quantum)}, \tag{44}$$

and the quantum comparison cell is the super-operator commutator

$$\mathcal{C}_Q[\rho] = -i[H, \mathcal{D}[\rho]] + i\,\mathcal{D}[[H, \rho]], \qquad \mathcal{C}_Q = [\mathcal{H}, \mathcal{D}]. \tag{45}$$

This analogy is not merely formal. Fix a faithful stationary state σ and equip $B(\mathbb{C}^n)$ with the *KMS inner product*

$$\langle A, B\rangle_\sigma := \mathrm{Tr}\big[\sigma^{1/2} A^\dagger \sigma^{1/2} B\big], \tag{46}$$

a real Hilbert-space structure on the Hermitian matrices, and define the superoperator path-reversal involution $\Theta(\mathcal{L}) := -\mathcal{L}^{\dagger_\sigma}$ exactly as in Definition 5.2. (The KMS weighting, rather than the GNS one, is the choice for which thermal jump dissipators are symmetric; see [19] for the comparison of the two conventions.)

Lemma 8.2 (Transfer of the structural framework)**.** *All structural results of §5.6 and §6.6—the Cartan symmetric-pair relations and cell grading (Lemma 5.3, Corollary 5.1), the ambidextrous adjunction (Theorem 5.1), the uniqueness of the natural splitting (Theorem 5.2), the global Cartan factorization (Theorem 5.3), and the exact interchange cocycle identities (Theorem 6.2) with their finite-size bounds (Theorem 6.3)—hold verbatim for Lindblad dynamics on* $\big(B(\mathbb{C}^n), \langle\cdot,\cdot\rangle_\sigma\big)$ *with* $\Theta(\mathcal{L}) = -\mathcal{L}^{\dagger_\sigma}$.

Proof. By inspection: every proof cited uses only (a) the structure of a real Hilbert space with its adjoint, and (b) the one-parameter group/semigroup laws. Both are available on the superoperator level; no property specific to commutative (classical) state spaces enters. □

Lemma 8.3 (Quantum detailed balance and the position of $\mathcal{H}$)**.** *Let* $[H, \sigma] = 0$. *Then* $\mathcal{H} = -i[H, \cdot]$ *is anti-self-adjoint for (46):* $\Theta(\mathcal{H}) = \mathcal{H} \in \mathfrak{k}$. *If moreover* $\mathcal{D}$ *is KMS-self-adjoint (*$\Theta(\mathcal{D}) = -\mathcal{D}$, *i.e.* $\mathcal{D} \in \mathfrak{p}$*)—the quantum detailed balance condition of Alicki and Kossakowski–Frigerio–Gorini–Verri [17, 18] in its KMS form—then the* Θ*-split of* $\mathcal{L}$ *is exactly* $\mathcal{L}_A = \mathcal{H}$, $\mathcal{L}_S = \mathcal{D}$. In contrast to the classical case, the antisymmetric part does not vanish at equilibrium*: quantum detailed balance leaves a coherent rotation* $\mathcal{H}$ *in* $\mathfrak{k}$.

Proof. With $[H, \sigma] = 0$, $\sigma^{1/2}$ commutes with H and cyclicity of the trace gives, for Hermitian A, B,

$$\langle A, \mathcal{H}B\rangle_\sigma = -i\,\mathrm{Tr}[\sigma^{1/2} A \sigma^{1/2}(HB - BH)] = -\big(-i\,\mathrm{Tr}[\sigma^{1/2}(HA - AH)\sigma^{1/2} B]\big) = -\langle \mathcal{H}A, B\rangle_\sigma,$$

so $\mathcal{H}^{\dagger_\sigma} = -\mathcal{H}$, i.e. $\Theta(\mathcal{H}) = \mathcal{H}$. The remaining statements are immediate from the definitions. (For the thermal jump dissipators used below, KMS-self-adjointness of $\mathcal{D}$ holds and is verified to machine precision.) □

Lemma 8.4 (Classical limit: commuting jump operators)**.** *If* $[H, L_k] = 0$ *for all* k, *then* $[\mathcal{H}, \mathcal{D}] = 0$, *i.e.* $\mathcal{C}_Q = 0$.

Proof. Since H is Hermitian, $[H, L_k] = 0$ implies $[H, L_k^\dagger] = -([H, L_k])^\dagger = 0$, hence $[H, L_k^\dagger L_k] = 0$. Then for every ρ: $[H, L_k \rho L_k^\dagger] = L_k[H, \rho]L_k^\dagger$ and $[H, \{L_k^\dagger L_k, \rho\}] = \{L_k^\dagger L_k, [H, \rho]\}$, so $[H, \mathcal{D}[\rho]] = \mathcal{D}[[H, \rho]]$, which is $\mathcal{C}_Q[\rho] = 0$ by (45). (For jump dissipators, Bohr-diagonal H—the secular condition—also gives $[\mathcal{H}, \mathcal{D}] = 0$ even when $[H, L_k] \neq 0$; verified numerically below.) □

Coherence parity. The quantum setting carries a second grading absent classically. Fix the eigenbasis of σ and let $\mathcal{C}(A) := \bar{A}$ (entrywise complex conjugation) on the real Hilbert space V of Hermitian matrices; $\mathcal{C}$ is an involutive isometry of (46). Call a superoperator X $\mathcal{C}$*-even* if $\mathcal{C}X\mathcal{C} = X$ and $\mathcal{C}$*-odd* if $\mathcal{C}X\mathcal{C} = -X$. For real H and real jump rates, $\mathcal{H} = -i[H, \cdot]$ is $\mathcal{C}$-odd while $\mathcal{D}$ is $\mathcal{C}$-even: coherent rotation and dissipation have opposite coherence parity.

Lemma 8.5 (Parity orthogonality)**.** *If X is $\mathcal{C}$-odd and Y is $\mathcal{C}$-even, then $\langle X, Y\rangle_{\mathrm{HS}} := \mathrm{Tr}[X^T Y] = 0$ (trace pairing of superoperators on V).*

Proof. Conjugation by the isometry $\mathcal{C}$ preserves the trace pairing: $\mathrm{Tr}[X^T Y] = \mathrm{Tr}[(\mathcal{C}X\mathcal{C})^T(\mathcal{C}Y\mathcal{C})] = \mathrm{Tr}[(-X)^T Y] = -\mathrm{Tr}[X^T Y]$. □

8.4 Quantum order separation

We can now prove the quantum analogue of Proposition 6.1—and it is *not* a verbatim copy of the classical statement: coherence changes the exponents.

Consider a smooth family $\mathcal{L}(\varepsilon) = \mathcal{H} + \mathcal{D}(\varepsilon)$ of Lindbladians satisfying:

(H1) *Equilibrium at $\varepsilon = 0$*: $\mathcal{L}(0)$ has faithful stationary state σ with $[H, \sigma] = 0$ and $\mathcal{D}(0)$ KMS-self-adjoint (quantum detailed balance, Lemma 8.3).

(H2) *Reversal-antisymmetric driving*: $\mathcal{D}(-\varepsilon) = \mathcal{D}(\varepsilon)^{\dagger_\sigma}$ (the quantum analogue of tilting the cycle affinity by $e^{\pm\varepsilon/2}$).

(H3) *$\mathcal{C}$-covariance*: in the σ-eigenbasis, H is real and $\mathcal{C}\mathcal{D}(\varepsilon)\mathcal{C} = \mathcal{D}(\varepsilon)$.

Write $\mathcal{K}_H, \mathcal{K}_D(\varepsilon)$ for the representations of $\mathcal{H}, \mathcal{D}(\varepsilon)$ in the fixed equilibrium KMS space, split as in Lemma 8.2, and let $\mathcal{A}(\chi_Q)(\varepsilon) = \frac{1}{2}\|[\mathcal{K}_A, \mathcal{K}_S](\varepsilon)\|_F$. Define the housekeeping entropy production through the quantum-jump unraveling [22]: for jump operators $L_{ij} = \sqrt{k_{ij}}\,|i\rangle\langle j|$,

$$\sigma^Q_{\mathrm{hk}} = \frac{1}{2}\sum_{i\neq j}(J_{ij} - J_{ji})\ln\frac{J_{ij}}{J_{ji}}, \qquad J_{ij} = k_{ij}\,\rho^{\mathrm{ss}}_{jj}. \tag{47}$$

Theorem 8.2: Quantum order separation

Under (H1)–(H3), with the NESS depending smoothly on ε (simple zero eigenvalue):

(i) **Coherent cost at equilibrium**: $\mathcal{A}(\chi_Q)(0) = \frac{1}{2}\|[\mathcal{K}_H, \mathcal{K}_D]\|_F \geq 0$, vanishing under the secular condition (in particular if $[H, L_k] = 0$, Lemma 8.4); generically $\mathcal{A}(\chi_Q)(0) > 0$.

(ii) **Coherent case** ($\mathcal{A}(\chi_Q)(0) > 0$): the first-order response vanishes by coherence parity,
$$\mathcal{A}(\chi_Q)(\varepsilon) = \mathcal{A}(\chi_Q)(0) + O(\varepsilon^2).$$

(iii) **Secular case** ($\mathcal{A}(\chi_Q)(0) = 0$): the classical exponent is recovered, $\mathcal{A}(\chi_Q)(\varepsilon) = c_1|\varepsilon| + O(\varepsilon^2)$ with $c_1 = \frac{1}{2}\|[\dot{\mathcal{K}}_D, \mathcal{K}_D]\|_F$.

(iv) $\sigma^Q_{\mathrm{hk}}(\varepsilon) = c_2\varepsilon^2 + O(\varepsilon^3)$.

(v) Hence near equilibrium $\mathcal{A}(\chi_Q)/\sigma^Q_{\mathrm{hk}} \to \infty$ like ε^{-2} in the coherent case and ε^{-1} in the secular case: *quantum coherence shifts the order separation from $(1, 2)$ to $(0, 2)$.*

Proof. *(i)* is Lemma 8.3 plus Lemma 8.4.

Expansions. By (H2), in the fixed equilibrium KMS representation $\mathcal{K}_D(\varepsilon)^\dagger = \mathcal{K}_D(-\varepsilon)$, so the symmetric part $\mathcal{K}_S(\varepsilon) = \frac{1}{2}(\mathcal{K}_D(\varepsilon) + \mathcal{K}_D(-\varepsilon))$ is an *even* function of ε, $\mathcal{K}_S = \mathcal{K}_D(0) + O(\varepsilon^2)$, while the antisymmetric part is $\mathcal{K}_A(\varepsilon) = \mathcal{K}_H + \varepsilon\dot{\mathcal{K}}_D + O(\varepsilon^3)$. Hence

$$C(\varepsilon) := [\mathcal{K}_A, \mathcal{K}_S](\varepsilon) = C_0 + \varepsilon C_1 + O(\varepsilon^2), \qquad C_0 = [\mathcal{K}_H, \mathcal{K}_D(0)], \quad C_1 = [\dot{\mathcal{K}}_D, \mathcal{K}_D(0)].$$

(ii) By (H3) and Lemma 8.3, $\mathcal{K}_H$ is $\mathcal{C}$-odd while $\mathcal{K}_D(0)$ and $\dot{\mathcal{K}}_D$ are $\mathcal{C}$-even (the equilibrium KMS conjugation is by real matrices, so it preserves parity). Conjugation by $\mathcal{C}$ is multiplicative on commutators, so C_0 is $\mathcal{C}$-odd (one coherent letter) and C_1 is $\mathcal{C}$-even (none). By Lemma 8.5, $\langle C_0, C_1\rangle = 0$, so

$$\|C(\varepsilon)\|_F^2 = \|C_0\|_F^2 + 2\varepsilon\langle C_0, C_1\rangle + O(\varepsilon^2) = \|C_0\|_F^2 + O(\varepsilon^2),$$

and taking square roots at $\|C_0\|_F > 0$ gives the claim.

(iii) With $C_0 = 0$, $C(\varepsilon) = \varepsilon C_1 + O(\varepsilon^2)$ and $\mathcal{A} = \frac{1}{2}|\varepsilon|\,\|C_1\|_F + O(\varepsilon^2)$.

(iv) Identical to the classical argument of Proposition 6.1: at $\varepsilon = 0$ the currents balance, $J_{ij} = J_{ji}$ (detailed balance of the populations under (H1)); smoothness of the NESS gives $J_{ij} - J_{ji} = O(\varepsilon)$ and the log-affinities are $O(\varepsilon)$, so the sum is $O(\varepsilon^2)$ with c_2 the Onsager coefficient. *(v)* follows from (i)–(iv). □

Numerical verification. A qutrit family realizing (H1)–(H3): thermal ring jumps with potential $\phi = (0, 0, 1)$, asymmetric base weights ($w_{01} = 1.3$, others 1), cycle tilt $e^{\pm\varepsilon/2}$, and coherent coupling $H = \Omega(|0\rangle\langle 1| + |1\rangle\langle 0|)$, $\Omega = 0.8$, on the degenerate pair (so $[H, \sigma] = 0$). Numerically:

quantity	value / exponent	theorem
KMS-Hermiticity of dissipator	4.5×10^{-16}	(H1)
coherent cost $\mathcal{A}(\chi_Q)(0)$	1.4708	(i)
exponent of $\lvert\mathcal{A}(\varepsilon) - \mathcal{A}(0)\rvert$	2.0002	(ii)
exponent of σ_{hk}^Q	2.0001	(iv)
secular case (H Bohr-diagonal): $\mathcal{A}(0)$	3×10^{-15}	(i)
secular case: exponent of $\mathcal{A}(\varepsilon)$	1.0004	(iii)

Interpretation. The equilibrium mixing cost $\mathcal{A}(\chi_Q)(0) = \frac{1}{2}\|[\mathcal{H}, \mathcal{D}]\|$ is precisely the comparison-cell reading of Hasegawa's coherence correction $C(\tau) > 0$: a quantum system pays a categorical mixing cost *even at detailed balance*, because quantum detailed balance leaves the coherent rotation $\mathcal{H}$ in the reversible direction. Driving adds to this only at second order. The bigrading behind this—each quantum cell letter carries (reversal type, coherence parity): $\mathcal{H} = (\mathfrak{k}, \text{odd})$, $\mathcal{D} = (\mathfrak{p}, \text{even})$, driving $\dot{\mathcal{D}} = (\mathfrak{k}, \text{even})$—shows that coherent rotation and dissipative driving share their reversal type but differ in coherence parity, and only matching-parity sectors couple at first order.

8.5 Main theorem: unified TUR

Theorem 8.3: Unified TUR

For an open quantum Markov system $(H, \{L_k\})$:

(i) **Classical limit** ($[H, L_k] = 0$, secular approximation): $\mathcal{C}_Q = 0$, $B(\tau) = A_{\rm cl}(\tau)$, and the TUR reduces to the classical bound.

(ii) **General case**: $C(\tau) = \int_0^\tau \mathrm{Tr}[\rho(t)\, \mathcal{C}_Q^* \mathcal{C}_Q[\mathbb{1}]]\, dt \sim \int_0^\tau F_Q^{\rm coh}[\rho(t)]\, dt$, where the "$\sim$" is controlled by Theorem 8.5: the two sides differ by $O(\Lambda\, \delta_{\rm sec}\, \tau^2)$ with $\delta_{\rm sec}$ the generator-level secular error.

(iii) **Comparison-cell identity**: $C(\tau) = \mathcal{A}(\square_Q) - \mathcal{A}(\Phi(\square_Q))$, where $\Phi : \mathbb{T}\mathrm{herm}_Q \to \mathbb{T}\mathrm{herm}_{\rm cl}$ is the coherent–incoherent correspondence functor of §8.7.

Proof status. Claim (i) is proved: Lemma 8.4 gives $\mathcal{C}_Q = 0$ from $[H, L_k] = 0$, whence $C(\tau) = 0$ and $B(\tau) = A_{\rm cl}(\tau)$; the classical bound then follows with Theorem 8.1 and the following quality-factor bound, which is fully proved:

Theorem 8.4: $\mathcal{Q}_{\min} \geq 1$

If $A_{\rm cl} > 0$, $C \geq 0$, and $C \leq A_{\rm cl}$, then

$$\mathcal{Q}_{\min} := \frac{2A_{\rm cl}}{A_{\rm cl} + C} \;\geq\; 1. \tag{48}$$

Proof. $C \leq A_{\rm cl}$ gives $2A_{\rm cl} \geq A_{\rm cl} + C > 0$; dividing by $A_{\rm cl} + C$ yields the claim. □

Claim (iii) is definitionally true once $\mathcal{A}$ and Φ are defined as stated. Claim (ii) is quantified by the following subsection.

8.6 Secular error propagation: quantifying claim (ii)

The validity of $C(\tau) \sim \int F_Q^{\rm coh}\, dt$ rests on replacing the exact reduced dynamics by a secular (Davies-type) Lindbladian. The *model-independent* part of the error analysis is a clean perturbation statement, which we prove; the *model-dependent* part—the size of the generator-level error—is the content of the rigorous weak-coupling literature [20, 21] and enters only as a single parameter.

Lemma 8.6 (Semigroup perturbation)**.** *Let $\mathcal{L}, \mathcal{L}'$ generate contraction semigroups on a Banach space (for CPTP semigroups: trace-norm contractions) with $\|\mathcal{L} - \mathcal{L}'\| \leq \delta$. Then $\|e^{t\mathcal{L}} - e^{t\mathcal{L}'}\| \leq \delta\, t$ for all $t \geq 0$.*

Proof. Duhamel's formula gives $e^{t\mathcal{L}} - e^{t\mathcal{L}'} = \int_0^t e^{(t-s)\mathcal{L}}(\mathcal{L} - \mathcal{L}')e^{s\mathcal{L}'}\, ds$; both outer factors have norm ≤ 1, so the integrand has norm $\leq \delta$ and the integral is bounded by δt. □

Theorem 8.5: Secular error propagation

Let $\rho(t) = e^{t\mathcal{L}}\rho_0$ and $\rho^{\rm sec}(t) = e^{t\mathcal{L}_{\rm sec}}\rho_0$ with $\delta_{\rm sec} := \|\mathcal{L} - \mathcal{L}_{\rm sec}\|$, and let g be Lipschitz with constant Λ on a convex set containing both trajectories. Then

$$\left| \int_0^\tau g(\rho(t))\, dt - \int_0^\tau g(\rho^{\rm sec}(t))\, dt \right| \;\leq\; \frac{\Lambda\, \delta_{\rm sec}\, \tau^2}{2}. \tag{49}$$

In particular: (a) for the linear integrand of $C(\tau)$, $g(\rho) = \mathrm{Tr}[\rho\,\mathcal{X}]$, one may take $\Lambda = \|\mathcal{X}\|_\infty$; (b) for $g = F_Q^{\mathrm{coh}}[\cdot]$, the bound applies on any trajectory tube where the state spectrum is bounded below by $\lambda_0 > 0$, with $\Lambda = O(\lambda_0^{-2})$ from the eigenvalue denominators in (41). The validity condition for claim (ii) is therefore $\Lambda\,\delta_{\mathrm{sec}}\,\tau^2 \ll C(\tau)$, with the entire secular-approximation dependence isolated in the single parameter δ_{sec}, which vanishes in the Davies weak-coupling limit [20] and admits quantitative bounds [21].

Proof. By Lemma 8.6, $\|\rho(t) - \rho^{\mathrm{sec}}(t)\| \le \delta_{\mathrm{sec}}\,t$, so

$$\left|\int_0^\tau \big(g(\rho) - g(\rho^{\mathrm{sec}})\big)dt\right| \le \Lambda \int_0^\tau \|\rho(t) - \rho^{\mathrm{sec}}(t)\|\,dt \le \Lambda\,\delta_{\mathrm{sec}} \int_0^\tau t\,dt = \frac{\Lambda\,\delta_{\mathrm{sec}}\,\tau^2}{2}. \qquad \square$$

This settles parts (a) and (b) of the proof audit (problem P6 of the resolution record, §9.2): the error order is explicit and quadratic in τ, and the secular dependence enters through δ_{sec} alone; what remains model-dependent is the standard evaluation of δ_{sec} for a given bath spectral density (open problem Q4).

8.7 Coherent–incoherent correspondence functor Φ

$\Phi : \mathbb{T}\mathrm{herm}_Q \to \mathbb{T}\mathrm{herm}_{\mathrm{cl}}$ acts by $\Phi(\rho) = \rho_{\mathrm{diag}}$ on objects, $\Phi(h_t^Q) = \mathrm{id}$ on horizontal morphisms (coherent rotations are invisible classically), and $\Phi(v_s^Q) = v_s^{\mathrm{cl}}$ on vertical morphisms (the dissipator descends to the classical kernel of populations).

8.8 Explicit calculation: field-driven two-level system

For a two-level system with $H = (\Omega/2)\sigma_x$ coupled to a thermal bath: $C(\tau) \approx \tau\Omega^2|\rho_{12}|^2 f(\gamma, T)$ and $\mathcal{Q}_{\min} \approx 2/(1 + C/A_{\mathrm{cl}})$, consistent with the reported value $\mathcal{Q}_{\min}^{\mathrm{TLS}} \approx 1.25$ [25].

8.9 New predictions

A: $C_{A\otimes B} = C_A + C_B + C_{\mathrm{int}}$ (non-additivity of coherence corrections). **B**: $C(\tau)/A_{\mathrm{cl}}(\tau) = g(\Omega^2/\gamma^2,\, T/\Omega)$ (universal scaling function). **C**: $C \le A_{\mathrm{cl}}$, hence $\mathcal{Q}_{\min} \ge 1$ (Theorem 8.4).

9 Discussion and Outlook

9.1 What has been achieved

1. **Conceptual foundation**: The Duality Proposition avoids both naive independence and strong reductionism.

2. **Mathematical structure**: $\mathbb{T}\mathrm{herm}$ is a lax double category; the time-like and entropy-like directions are the isotypic components of the path-reversal involution, characterized—for arbitrary probability spaces and with no reference to entropy—by the ambidextrous adjunction of Theorem 5.1, unique among natural splittings (Theorem 5.2), and globally jointly exhaustive via the Cartan factorization (Theorem 5.3). This breaks the circularity of naive approaches, and the Cartan relations (6) assign every comparison cell of the BCH hierarchy a definite reversible/dissipative

type (Corollary 5.1). The GNS construction (Appendix A) resolves the type-mismatch of the naive Helmholtz split (Theorem 6.1). The lax coherence of the comparison cell holds *exactly and at all orders* as a pair of interchange cocycle identities (Theorem 6.2), with the pentagon and hexagon as linearizations, exact horizontal subadditivity of the invariant cost (Lemma 6.11), and two-sided finite-size bounds for the deviation from the area law (Theorem 6.3).

3. **Stochastic thermodynamics**: The Order Separation Proposition establishes universal exponents $\mathcal{A}(\chi) \sim \varepsilon^1$ and $\sigma_{\mathrm{hk}} \sim \varepsilon^2$ on the generic locus, proved analytically and confirmed computationally. The zero sets are nested, $\{H_A = 0\} \subsetneq \{[H_A, H_S] = 0\}$ (Theorem 6.4), with circulant chains as the strict stratum. $\mathcal{A}(\chi)$ is the first-order (linear-response) order parameter of non-equilibrium physics.

4. **Exact results**: closed forms for the two-bath model (Theorem 6.5) and the counterexample pair CE-1/CE-2; the channel-resolved (bath-modulated) lift of §6.10 restores $H_A \neq 0$ for the driven two-bath system, identifies the physical driving δ with the cycle affinity (Lemma 6.14), and shows that CE-1/CE-2 are artifacts of the channel-summed description.

5. **Quantum thermodynamics**: the structural framework transfers verbatim to Lindblad dynamics via the KMS-superoperator involution (Lemma 8.2); quantum order separation is *proved* (Theorem 8.2), with coherence shifting the exponents from $(1, 2)$ to $(0, 2)$—the equilibrium mixing cost $\mathcal{A}(\chi_Q)(0) = \frac{1}{2}\|[\mathcal{H}, \mathcal{D}]\|$ being the comparison-cell reading of the TUR coherence correction; the two-level quantum Cramér–Rao bound and $\mathcal{Q}_{\min} \geq 1$ are proved, the classical limit of the unified TUR is proved (Lemma 8.4), and the secular error is quantified (Theorem 8.5).

9.2 Resolution record

The development of this framework passed through six technical problems, P1–P6. All six have been settled—five fully, one partially—and their resolutions are integrated into the body of the paper. The record:

Problem	Status	Resolution
P1: full lax coherence tower	Resolved	Exact group-level interchange cocycle identities, all BCH orders, no Poincaré–Birkhoff–Witt machinery (Theorem 6.2); pentagon and hexagon are the linearizations (Corollary 6.1).
P2: interpretation of c_1/c_2	Partial	Exact spectral representation of c_1 as spectral-gap-weighted coupling (Corollary 6.2); the ratio measures driving–spectrum alignment. Classification question remains (Q1).
P3: adjoint existence in full generality	Resolved	The two directions are the isotypic components of the path-reversal involution: ambidextrous adjunction (Theorem 5.1), uniqueness of the natural splitting (Theorem 5.2), global Cartan factorization (Theorem 5.3); valid for arbitrary probability spaces, with unbounded generators under the weak sector condition (Remark 5.3).
P4: exact area law at finite (t, s)	Resolved (negatively)	The exact law fails at finite size, with two-sided bound $C\,ts(t+s)$, numerically sharp (Theorem 6.3); exact horizontal subadditivity of the invariant cost (Lemma 6.11).
P5: channel-resolved two-bath extension	Resolved	Bath-modulated lift restores $H_A \neq 0$; cycle affinity $= \delta$ exactly; detailed balance iff $\delta = 0$ (Lemma 6.14); CE-1/CE-2 are channel-summed artifacts (Remark 6.2).
P6: proof audit of the quantum TUR section	Resolved	Classical limit proved (Lemma 8.4); quantum order separation proved with exponents shifted by coherence from $(1, 2)$ to $(0, 2)$ (Theorem 8.2); secular error quantified as $O(\Lambda\,\delta_{\mathrm{sec}}\,\tau^2)$ (Theorem 8.5).

All resolutions are backed by numerical computation (cocycle identities and Cartan/parity gradings to machine precision; scaling exponents to 4 significant digits).

9.3 Remaining open problems

What follows are the genuine residuals, renumbered.

(Q1) Universality classes of c_1/c_2. The ratio measures the alignment between the driving direction and the relaxation spectrum (Corollary 6.2); classify the families of chains on which it is constant, and determine whether the alignment admits a variational characterization.

(Q2) Beyond the sector condition. Extend the involution framework of §5.6 to genuinely non-sectorial unbounded generators (antisymmetric part unbounded relative to the symmetric part), where the form-level splitting of Remark 5.3 is unavailable; and settle the uniqueness of the natural splitting within the smaller class of system intertwiners (Remark 5.4).

(Q3) Global cost. Identify the preferred globally defined cost for large (t, s), where the principal logarithm of Definition 6.1 may leave its domain; the invariant distance $\widehat{\mathcal{A}}$ of Lemma 6.11 is globally defined and is the natural candidate, but its exact relation to thermodynamic quantities at finite size is open.

(Q4) Quantitative secular input. Evaluate $\delta_{\mathrm{sec}} = \|\mathcal{L} - \mathcal{L}_{\mathrm{sec}}\|$ for specific bath spectral densities (the model-dependent input to Theorem 8.5), and control the Lipschitz constant of F_Q^{coh} near the boundary of the state space ($\lambda_0 \to 0$).

Relativistic open problems (observer-indexed stacks, the Unruh response computation, the normalization of horizon entropy, Bisognano–Wichmann, the generalized second law) are catalogued in Part II [26].

10 Conclusion

We have developed a unified framework for time–entropy duality through the lax double category $\mathbb{T}\mathrm{herm}$. Time and entropy both arise from the symmetry structure of the same path-space action functional.

The GNS decomposition $H = H_S + H_A$ provides the type-coherent foundation: horizontal morphisms are orthogonal maps preserving D_{χ^2} exactly; vertical morphisms are KL-contractive kernels. The comparison cells—measured by $[H_A, H_S]$—emerge as the categorical mixing cost, and their coherence is not merely perturbative: the exact interchange cocycle identities (Theorem 6.2) hold at all orders, with the reversible direction whiskering isometrically and the dissipative direction amplifying the cost by at most the relaxation factor—a cost-level signature of the time/entropy asymmetry. The Order Separation Proposition establishes that on the generic locus $\mathcal{A}(\chi)$ and σ_{hk} vanish at different orders (ε^1 and ε^2), universally across chain architectures, proved analytically and confirmed computationally; their zero sets are nested rather than equal, with circulant chains witnessing the strict inclusion. This order separation is the mathematical rationale for the primacy of $\mathcal{A}(\chi)$: it is the first-order (linear-response) order parameter of non-equilibrium physics.

The framework transfers verbatim to quantum Markov dynamics through the KMS-superoperator involution, and there the order separation is *modified* in a characteristic way: quantum detailed balance retains the coherent rotation in the reversible direction, the mixing cost acquires a non-zero equilibrium value—the comparison-cell reading of the TUR coherence correction—and coherence parity suppresses the first-order response, shifting the exponents from $(1, 2)$ to $(0, 2)$ (Theorem 8.2). Coherence, in this precise sense, is a zeroth-order thermodynamic resource.

The extension of these structures to relativistic observers, horizons, and black-hole entropy is developed in Part II [26].

References

[1] R. Landauer, *IBM J. Res. Dev.* **5**, 183 (1961).

[2] C. H. Bennett, *IBM J. Res. Dev.* **17**, 525 (1973).

[3] A. Bérut et al., *Nature* **483**, 187 (2012).

[4] S.-i. Sasa and Y. Yokokura, *Phys. Rev. Lett.* **116**, 140601 (2016).

[5] P. Erker et al., *Phys. Rev. X* **7**, 031022 (2017).

[6] A. N. Pearson et al., *Phys. Rev. X* **11**, 021029 (2021).

[7] A. Beyen and C. Maes, *J. Non-Equilib. Thermodyn.* **50**, 241 (2025).

[8] A. Beyen and C. Maes, *Continuum Mech. Thermodyn.* **37**, 84 (2025).

[9] Y. Minami and S.-i. Sasa, *J. Stat. Mech.* **2020**, 013213 (2020).

[10] A. Connes and C. Rovelli, *Class. Quantum Grav.* **11**, 2899 (1994).

[11] T. Jacobson, *Phys. Rev. Lett.* **75**, 1260 (1995).

[12] S. Ryu and T. Takayanagi, *Phys. Rev. Lett.* **96**, 181602 (2006).

[13] M. Grandis and R. Paré, *Cah. Topol. Géom. Différ. Catég.* **40**, 162 (1999).

[14] C. Ehresmann, *Ann. Sci. École Norm. Sup.* **80**, 349 (1963).

[15] Z.-M. Ma and M. Röckner, *Introduction to the Theory of (Non-Symmetric) Dirichlet Forms*, Springer (1992).

[16] S. Helgason, *Differential Geometry, Lie Groups, and Symmetric Spaces*, Academic Press (1978).

[17] R. Alicki, *Rep. Math. Phys.* **10**, 249 (1976).

[18] A. Kossakowski, A. Frigerio, V. Gorini, M. Verri, *Commun. Math. Phys.* **57**, 97 (1977).

[19] F. Fagnola and V. Umanità, *Infin. Dimens. Anal. Quantum Probab. Relat. Top.* **13**, 335 (2010).

[20] E. B. Davies, *Commun. Math. Phys.* **39**, 91 (1974).

[21] E. Mozgunov and D. Lidar, *Quantum* **4**, 227 (2020).

[22] J. M. Horowitz and J. M. R. Parrondo, *New J. Phys.* **15**, 085028 (2013).

[23] Y. Hasegawa, *Phys. Rev. Lett.* **126**, 010602 (2021).

[24] T. Nishiyama and Y. Hasegawa, arXiv:2505.09973 (2025).

[25] S. Cho et al., *Phys. Rev. Research* (2025).

[26] *Time–Entropy Duality and the Lax Double Category of Thermodynamics II: Relativistic Extension, Horizons, and Black-Hole Entropy*, companion paper (2026).

A GNS Transformation and the D1–D2 Construction

A.1 The ψ-space transformation

Let π be the stationary distribution of L_{fwd} and $\Pi = \text{diag}(\pi)$. Define:

$$\psi := \Pi^{-1/2}\rho, \quad \rho = \Pi^{1/2}\psi, \quad H := -\Pi^{-1/2}L_{\text{fwd}}\Pi^{1/2} = H_S + H_A. \tag{50}$$

A.2 Proof of $H_A\sqrt{\pi} = H_S\sqrt{\pi} = 0$

$H\sqrt{\pi} = -\Pi^{-1/2}L_{\text{fwd}}\,\pi = 0$ (stationarity, Lemma 6.2). $H^T = -\Pi^{1/2}L_{\text{fwd}}^T\Pi^{-1/2}$, so $H^T\sqrt{\pi} = -\Pi^{1/2}L_{\text{fwd}}^T\mathbf{1} = 0$ (column sums zero, Lemma 6.1). Hence $H_S\sqrt{\pi} = \frac{1}{2}(H + H^T)\sqrt{\pi} = 0$ and $H_A\sqrt{\pi} = \frac{1}{2}(H - H^T)\sqrt{\pi} = 0$.

A.3 D1–D2 construction

Push-forward on distributions: $\rho' := h_t^*\rho = \Pi^{1/2}e^{-tH_A}\Pi^{-1/2}\rho$.

Normalization: $\sum_x \rho'_x = \sqrt{\pi}^T e^{-tH_A}\psi$. Since $H_A^T\sqrt{\pi} = -H_A\sqrt{\pi} = 0$ (antisymmetry plus the kernel lemma), the row vector $\sqrt{\pi}^T$ is a left eigenvector of e^{-tH_A} with eigenvalue 1. Hence $\sqrt{\pi}^T e^{-tH_A}\psi = \sqrt{\pi}^T\psi = \sum_x \rho_x = 1$.

D_{χ^2} preservation: In ψ-coordinates, $D_{\chi^2}(\rho\|\pi) = \|\psi - \sqrt{\pi}\|_2^2$. Since $e^{-tH_A} \in O(n)$ and $e^{-tH_A}\sqrt{\pi} = \sqrt{\pi}$:

$$D_{\chi^2}(\rho'\|\pi) = \|e^{-tH_A}\psi - e^{-tH_A}\sqrt{\pi}\|_2^2 = \|e^{-tH_A}(\psi - \sqrt{\pi})\|_2^2 = \|\psi - \sqrt{\pi}\|_2^2 = D_{\chi^2}(\rho\|\pi). \tag{51}$$

A.4 Computational verification

The GNS split function implements the above. For any detailed-balance generator: $\|H_A\| \sim 10^{-16}$ (machine zero). D_{χ^2} preservation holds to machine precision across all test cases.

A.5 Relation to the Helmholtz formulation

$H_A = \Pi^{-1/2}L^{\text{antisym}}\Pi^{1/2}$ holds on the restricted domain where this is defined. However, $L^{\text{antisym}} = \frac{1}{2}(L_{\text{fwd}} - L_\pi^\dagger)$ mixes the forward-distribution and backward-observable pictures, making the difference ill-typed as an operator on a single space (Theorem 6.1). The GNS split works entirely in ψ-space, where a single Euclidean structure governs both parts.